Pricing-to-Market and the Failure of Absolute PPP

By George Alessandria and Joseph P. Kaboski*

Abstract: We show that deviations from the law of one price in tradable goods are an important source of violations of absolute PPP across countries. Using highly disaggregated export data, we document systematic international price discrimination: at the U.S. dock, U.S. exporters ship the same good to low-income countries at lower prices. This pricing-to-market is about twice as important as any local non-traded inputs, such as distribution costs, in explaining the differences in tradable prices across countries. We propose a model of consumer search that generates pricing-to-market. In this model, consumers in low-income countries have a comparative advantage in producing non-traded, non-market search activities and therefore are more price sensitive than consumers in high-income countries. We present cross-country time use evidence and evidence from U.S. export prices that are consistent with the model. (JEL E31, F12)

Keywords: PPP, Pricing-To-Market, Law of One Price.

[INSERT FIGURE 1 HERE]

Figure 1 plots income per capita against the consumer price level and draws a line with the estimated magnitude of this relation for the benchmark countries in the Penn World Tables. This picture raises two questions: First, why are there large differences in consumer price levels across countries? The theory of absolute Purchasing Power Parity (PPP) states that the same basket of goods should sell for the same price everywhere, yet, for instance, the price level in Mexico is 64

*Alessandria: Research Department, Federal Reserve Bank of Philadelphia, Ten Indepedence Mall, Philadelphia, PA 19106 (george.alessandria@phil.frb.org); Kaboski: Department of Economics, 434 Flanner Hall, University of Notre Dame, Notre Dame, IN 46556 (jkaboski@nd.edu). Alessandria thanks the National Science Foundation for financial support. Kaboski thanks the National Institute of Health and The Ohio State Univesity for financial support. This paper is a substantially revised version of the working paper, "Violating Purchasing Power Parity." The authors thank two anonymous referees, Andy Atkeson, Ariel Burstein, Mario Crucini, Marvin Goodfriend, Patrick Kehoe, Fabrizio Perri, Kim Ruhl, Bruce Weinberg, and Kei-Mu Yi for helpful comments. Virgiliu Midrigan and Jarcy Zee provided excellent research assistance. All remaining errors are our own. The views expressed here are those of the authors and do not reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

percent of the price level in the United States.[1] Second, why are price levels so strongly related to income per capita? A doubling of income per capita is associated with a 48 percent higher price level.[2] The conventional explanation for these two observations is the model of Roy F. Harrod (1933), Bela Balassa (1964), and Paul Samuelson (1964), the HBS model hereafter. In HBS, differences in price levels are driven solely by non-tradable goods, for which the law of one price (LOP) doesn't hold. Since the LOP holds for traded goods, international relative wages are determined by the large productivity differences in tradables. Large differences in wages lead to differences in the price of non-tradables, a sector in which productivity differences are much smaller across countries. The model therefore rests on the assumption that cross-country productivity differences are much smaller in non–tradables than in tradables, and that the LOP holds in tradables.[3]

[INSERT FIGURE 2 HERE]

There are two good reasons to doubt HBS as a full explanation of these observations. First, we see from Figure 2, which plots the price levels of tradable consumption goods against income per capita, that the LOP for tradables is clearly violated in the data. As in Figure 1, the relationship between prices and income is positive and significant, and the estimated elasticity in tradables (0.31) is nearly two-thirds of the overall elasticity (0.48).[4] Second, to explain the magnitude of the relation in Figure 1, the rise in relative productivity of tradables with income across countries (the cross-section) would have to be much bigger than what we observe within countries (the time series).

The aim of this paper is to evaluate the role of systematic price discrimination across coun-

[1] Kenneth Rogoff (1996) provides a review of PPP.

[2] We focus here on consumer prices, which is commonly the focus in the literature on PPP. Similar relationships hold when looking at all final goods, however, with elatisticities of 0.43 overall and 0.26 for tradables.

[3] A few existing theories present alternatives to HBS. Irving Kravis and Robert E. Lipsey (1983) and Jagdish N. Bhagwati (1984) focus on differences in factor endowments. Staffan B. Linder (1961), Rudiger Dornbusch (1988), Peter Neary (1988) and Jeffrey H. Bergstrand (1991) focus on differences in preferences. All of these theories assume that the LOP holds for tradables.

[4] These deviations from PPP are quite persistent. Of the countries with price levels less than half the U.S. level in 1996, over 80 percent (26 out of 32 available in 1985) also had price levels less than one-half that of the U.S. in 1985. These countries also had very low income, with GDP per capita on average 14 percent of the U.S. level in 1996.

tries, what Paul Krugman (1987) calls pricing-to-market, in the pattern of tradable and overall price levels in Figures 1 and 2. The tradable price variation in Figure 2 could be interpreted as direct evidence of pricing-to-market, but a potential problem is that tradable prices are measured using final goods prices. Thus, differences in retail prices may be driven by non-traded components, such as transportation and distribution, instead of differences in the actual price of tradables earned by the producer. Additionally, despite the best efforts of statistical agencies to compare identical goods, there may be quality differences between goods in different countries. Ideally, to overcome such data concerns requires micro price data on identical products in multiple countries. Such data is generally only available for a few goods (see Ina Simonovska, 2008, for some clothing products in Europe) and is suggestive. Clear examples of a good where the exporter price discriminates based on the destination market are submission fees to AEA journals, including this journal. Submission fees vary from \$0 to \$100 (\$0 to \$200 for non-members) depending on the country where the submitter works. Similarly, membership fees vary positively with income with a significant elasticity 0.06.[5]

Evaluating the role of pricing-to-market in aggregate PPP requires broader data, however. We examine unit values of the universe of U.S. exports, data for which the above concerns are (partially) mitigated. In particular, the data are well-suited for isolating pricing-to-market because they are highly disaggregate, and they measure the export income received at the U.S. border before any local non-traded inputs are added. The data show that unit values are higher for exports to high-income countries. On average, the richest country in our data set pays 48 percent more per unit than the poorest country. These results appear robust to controls for quality differences and suggest price discrimination is common. The estimated elasticity of price with respect to GDP per capita is 0.22, indicating that about two-thirds of deviations in the LOP in final tradables

[5]We cannot estimate a similar pricing-to-market coefficient for submissions, since submissions from low income countries are free. However, restricting ourselves to countries with positive submission fees we find a pricing-to-market coefficient of 38 percent. Details of our analysis of AEA subscription and membership prices are given in the unpublished web appendix.

could be due to pure pricing-to-market. Thus, this data show that the tradable price vs. income per capita relationship from Figure 2 may primarily reflect pricing-to-market rather than the non-traded content of tradable goods.

What characteristics of consumers or markets might lead firms to price discriminate based on income? The second aim of this paper is to propose a strong candidate for a micro-level explanation for higher elasticities of demand from consumers in low-wage countries. In Section II, we develop a model based on consumer search frictions and international productivity differences. Search requires time, which consumers in high-productivity/high-wage countries value at a premium. They are therefore less willing to search and less price elastic shoppers. Firms take this into account and set relatively high prices (and therefore markups) when selling to high-wage countries. Thus, there is a tight (endogenous) link between the local wage and prices, both tradable prices *and* non-tradable prices. Our consumer search story parallels and complements the HBS story. That is, we too rely on small differences in the productivity of a non-traded good. The search friction story requires that productivity in shopping rises less rapidly with income than productivity in market production.

Our focus on search as the source of pricing-to-market is motivated by the substantial evidence that search matters for explaining how prices vary by income within countries. Time-use studies find that poorer consumers, and those with a lower opportunity cost of time, spend relatively more time shopping per purchase (David McKenzie and Ernesto Schargrodsky, 2005, and Mark Aguiar and Erik Hurst, 2007). Furthermore, these studies find that shopping time is negatively related to purchase price, a direct implication of the search model. This effect is potentially large within countries. For instance, Aguiar and Hurst find, using scanner data on prices paid and time-use data on shopping time, that a doubling of search effort lowers the expected purchase price by 7 to 10 percent. Our search model can generate a similar relationship within countries, but because in general equilibrium firm pricing responds to an elastic demand, we get much larger effects across countries. For this reason, search is a quantitatively important source of pricing-to-market. A key advantage of our approach is that it offers a simple, unified theory to explain price dispersion within and across countries.

In Section III, we conduct a quantitative analysis of the search model. We find the model can account for about half of the observed pricing-to-market relationship, and 51 percent of the PPP-income relationship, or about twice as much as the HBS model alone explains. The existence of pricing-to-market augments the HBS explanation in two ways. It helps reconcile the smaller observed differences in the relative price of tradables to non-tradables, and it also helps reconcile large differences in average price levels with the evidence that relative productivity in the tradable sector does not increase nearly as much with income (in the time series) as the HBS explanation would require (in the cross-section of countries). Finally, if our model captures the main source of pricing-to-market, then the difference between our model's quantitative results and our empirical findings using export data suggest that quality variation accounts for at most one-half of the price-income relationship in the export data.

Further corroborating evidence, which we review, supports the consumer search model as a strong candidate explanation for the pricing-to-market we observe. First, we show, using cross-country time-use studies, that the ratio of shopping time to work time increases substantially with income, which indicates that shopping productivity does not increase as rapidly as income (and overall productivity). Second, our U.S. export evidence shows that the opportunity cost of time (wage) is more robustly associated with prices than with income, and that pricing-to-market is strongest for consumer goods. Finally, quantitatively, numerical examples indicate that the search model can potentially generate pricing-to-market of the order observed in the U.S. export data.

In addition to contributing to the study of absolute PPP, this paper relates to two other literatures. Our emphasis on pricing-to-market in tradables as an important source of violations from absolute PPP is consistent with the prevailing view in the literature on relative PPP. Charles Engel (1999) and V. V. Chari, Patrick J. Kehoe and Ellen McGrattan (2002) show that deviations from the LOP in tradables account for nearly all of the fluctuations in real exchange rates among developed countries.[6] Theoretical explanations of this pricing-to-market take two forms. The first approach focuses on the role of sticky prices set in local currencies, while the second emphasizes that local market conditions differ across countries and time so that firms have incentive to systematically price discriminate internationally.[7] Since we are looking at absolute PPP and long term deviations, we follow the approach of focusing on local market conditions. This paper also relates to the literature on the role of relative prices and productivities in capital accumulation and growth. Jonathan Eaton and Samuel Kortum (2001) and Chang-Tai Hsieh and Peter J. Klenow (2007) demonstrate that relative price differences across countries are important in explaining cross-country variation in capital stocks and income levels. Hsieh and Klenow therefore argue that understanding the origins of these relative price and productivity differences is essential. We argue that pricing-to-market, and not only relative productivities, plays a role in the prices (of investment, for example) that countries face.

I Pricing-to-Market: Empirics and Importance

In this section, we document evidence, using highly disaggregated data on U.S. exports, that U.S. firms systematically price discriminate by the income of the destination market. We show that this price discrimination is not likely due to unobserved quality differences, and that it provides some evidence for a search model where the elasticity of demand varies with the opportunity cost of

[6] Patrick Asea and Enrique Mendoza (1994) find the HBS model can not explain real exchange rate and output fluctuations.

[7] Starting with Krugman (1987), the local market condition models have sought to explain differences in elasticities of demand across countries from first principles without resorting to differences in tastes. A variety of local market condition models exist and emphasize both supply considerations, such as differences in industry structure (Rudiger Dornbusch, 1987), and demand considerations, such as the decisions of firms to build market shares (George Alessandria, 2004).

search. Finally, using a modified version of Engel's (1999) decomposition of real exchange rates, we find that pricing-to-market accounts for 40 percent of the aggregate price-income relationship and the non-traded component of final goods accounts for 20 percent.

A. Export Data

The micro data we analyze, U.S. Exports Harmonized System data (see Robert Feenstra et al., 2002), have significant advantages over the aggregate data in identifying pricing-to-market in tradables.

First, the data are comprehensive of all U.S. domestic exports (excluding re-exports) and therefore include only tradable goods. We focus on consumer goods but also present evidence for a much broader range of tradables. We have annual data on the total value and quantity of all commodities exported by destination country. We link these data to income per capita data from the Penn World Tables 6.1 for the years 1989-2000 and to hourly (manufacturing) wage data available from the BLS. These wages are reported in nominal local currency, which we convert to international Geary-Khamis dollars using the PWT PPP price level.[8]

Given our emphasis on search and the opportunity cost of time, we focus on countries for which both hourly wage and income per capita data are available. These 28 countries include most long-term members of the OECD plus Hong Kong, Israel, South Korea, Mexico, Singapore, Sri Lanka, and Taiwan. Over the 12 years of annual data, we have 1.1 million good-year observations for these countries, constituting 78 percent of the value of U.S. exports.

The second crucial advantage of this export data is that they are collected "at the dock" of the U.S. That is, our export prices are based on free-alongside-ship values,[9] so they do not

[8] Converting to U.S. dollars using exchange rates produces nearly identical results, except that estimated elasticities of prices are smaller with respect to exchange rate-based income per capita, given their larger variance.

[9] The free-alongside-ship value is the selling price or cost if not sold, including inland freight, insurance, and other charges to the U.S. port of export, but excluding unconditional discounts and commissions. It is essentially the price received by the exporting country before shipment.

include transportation costs, tariffs, or distribution and retailing costs in the importing country. One complication, however, is that we do not directly observe prices. Instead, we calculate unit values from data on the total value and quantity sold. Numerous important studies of deviations of the LOP and pricing-to-market are based on unit values (see Peter Isard, 1977, Robert Feenstra, 1989, and Michael Knetter, 1993). Unit values have the advantage of providing a measure of destination-specific prices for a large number of products.

A common concern with unit values is that they may combine goods of different qualities. These concerns are mitigated somewhat since we are using quite disaggregate data. Indeed, we have 10,741 products classified using the 10-digit Harmonized System product codes. (The appendix lists 10 randomly selected goods as an example of the level of detail. A broader random sample is given in our unpublished web appendix.) Nonetheless, quality variation remains a concern, and we consider possible biases below.

B. Pricing-to-Market Evidence

For exposition, consider a monopolist selling an identical good in different markets (e.g., countries). Facing different demand in each market, the firm will, in general, charge price p equal to a market-varying markup μ over a common marginal cost c. Hence, while marginal costs and markups may vary across goods, i, and time, t, markups also vary across destination market, j:

$$\ln p_{ijt} = \ln c_{it} + \mu_{ijt}.$$

The purpose is to examine whether μ_{ijt}, the markup charged on good i at time t to destination country j, is related to the level of income per capita or wages of that country. We estimate the following regression equation:

$$(1) \quad \ln p_{ijt} = \alpha_{it} + \beta \ln y_{jt} + e_{ijt},$$

in which y_{jt} is a measure of destination country income (either GDP per capita, wage, or a vector of both). The intercepts, α_{it}, capture variation in $\ln c_{it}$. They are estimated as fixed effects for each good-year combination. We use the "within" estimator, so that the identification of β comes from variation in the income of destination countries within good-year cells.[10] We report White robust standard errors that allow for heteroskedasticity in e_{ijt}, and also allow for country-year clustering.[11]

Table 1 presents the estimated $\hat{\beta}$ coefficients on log income and/or log wages from these fixed-effect regressions. In these baseline estimates, we focus on consumer goods and automotives, whose effect on the cost of consumption goods is most direct, but we present results for other types of traded goods in the next two subsections. The "GDP per Capita Only" estimate, from a regression where log GDP per capita is the only regressor (in addition to the fixed effects), yields an elasticity estimate of 0.235 on PPP income/capita. The "Wage Only" estimates are slightly smaller at 0.209.

Both sets of estimates are highly significant, but when we include log wages and log GDP per capita together in the same regression, log wages wins the horse race hands down. These estimates are presented in the right-most column. The estimated coefficient on wages remains at nearly the same level (0.188) and is highly significant, while the GDP per capita coefficient estimate becomes much smaller (0.041) and insignificant.[12]

Based on the estimates in Table 1, the magnitude of the price-wage relationship is potentially large. In 2000, the difference in log wages between the richest and poorest countries in the data set

[10] Marginal costs are made both good- *and* year-specific to avoid problems with changing quality over time and issues of non-stationarity in income and prices.

[11] Clustering on good-year has a minor effect on standard errors with the nearly 40,000 good-year combinations.

[12] Including measures of trade costs such as distance, tariffs, and downstream distribution costs does not substantially alter our estimates.

(Germany and Sri Lanka, respectively) was 2.4 measured in PPP terms. Hence, the implied price differences in U.S. exports to these countries would be 50 percent.

C. Quality

We interpret the unit value-income relationship we observe as pricing-to-market. Although the data are extremely disaggregated, one might still suspect that the positive relationship uncovered is driven by unobserved quality variation.[13] That is, perhaps our measured relationship simply reflects (in part or in whole) a tendency for wealthy countries to import higher quality (and higher priced) goods within the 10-digit commodity categories. Indeed, variation in import prices at the same level of aggregation has been attributed to such quality variation (see Peter Schott, 2004, and Juan Hallak, 2006). Unlike imports though, the quality variation in our data is mitigated somewhat since exports are from a single source country. Nonetheless, definitively distinguishing between price discrimination or quality differences requires price data such as what we observe for AEA memberships and submissions. Lacking such data for a broad set of goods, here we propose a number of tests that control for quality variation somewhat, and the results suggest that some of the observed pricing-to-market relationship may indeed result from price discrimination.

Our first approach follows Schott (2004) by using variation in unit values of imports within an HS-10 category as a measure of the extent of quality variation within the category.[14] Specifically, we define quality variation at the HS 10 level (for a particular year) as the standard deviation of ln unit values (σ_j) across all sources. This follows the convention in the literature that attributes price variation to quality variation. This convention is clearly at odds with the within-country evidence

[13] In addition to quality as a potential non-PTM explanations, we also evaluated whether the relationship could be driven by transfer pricing among related party trade, and the evidence did not support this. We also considered whether pricing might be driven by intellectual property protection and local competition from pirated goods. Our results were robust to inclusion of indexes of intellectual property protection.

[14] We considered a variety of alternate approaches to controlling for quality. For instance, we dropped those product categories viewed as potentially most heterogeneous such as those commodities with descriptions containing words like "other," "not elsewhere specified or included," "NESOI," and "parts." The elasticity estimates are quite similar to the full sample. For instance our full sample estimate in Table 2 falls from 0.159 to 0.149.

of price variation unrelated to quality variation, however. Additionally, this measure is likely to overstate quality variation as prices may also be dispersed, if costs are dispersed and producers have market power. In Table 2, we report pricing-to-market for goods with low quality variation and high quality variation, where goods have been divided by their dispersion in import unit values relative to the median σ_j. Our measure of pricing-to-market indeed increases with quality variation from 0.128 to 0.183, but the increase is small relative to the increase in quality; in the import data, those goods with high quality variation have nearly 2.5 times the average quality-driven unit value variation compared to low quality variation goods (σ of 1.92 vs. 0.78). Moreover, looking at the 10 percent of goods with the least quality variation ($\sigma_j < 0.46$) we still find substantial pricing-to-market coefficient of 0.107.[15] Finally, we can also run a simple regression of prices on wages that includes an interaction term with our quality measure ($\sigma_j \cdot w_{ij}$) of

$$(2) \quad p_{ij} - \overline{p}_j = \beta_0 \left(w_{ij} - \overline{w}_i\right) + \beta_1 \sigma_j \cdot \left(w_{ij} - \overline{w}_i\right) + \varepsilon_{ij},$$

where $\overline{p}_j \left(\overline{w}_i\right)$ is the average price (wage) across destinations. The coefficient on wages, β_0, here measures pricing-to-market on goods with no quality variation. Controlling for quality in this way yields a pricing-to-market coefficient of 0.084 or more than half the overall sample of 0.159. Thus, using a conventional measure of quality variation we still find substantial pricing-to-market.

A second approach is to further disaggregate the data by port and month of export. The presumption here is that goods sold from the same port in the same month are more likely to be from the same supplier and hence more likely to be of similar quality. A key caveat to this analysis is that ports may differ in their mode of transport (air, ship, truck, rail) and destinations and so some of the variation in shipments by ports is economically meaningful. Additionally, we do not have a

[15] Goods with this dispersion in unit values ($\sigma_j < 0.46$) have substantially less systematic variation in unit values related to the real wage of the source country compared to all goods. For instance, regressing unit values on source real wages yields a coefficient of 7.8 percent for low dispersion goods and a coefficeint of 66 percent for all imports.

theory of the timing of shipments,[16] and so it is not clear one wants to distinguish between monthly and annual purchases, particularly if, say, low income countries are more sensitive to price variation over the year. Nonetheless, we run our same regressions on monthly unit values of exports in two years, 1998 and 1999, but now defining goods at the port-month level. We also include a control for the number of transactions (to pick up volume discounts) and share of exports by air. Table 3 presents this result, which yields an elasticity of 0.096. Our estimated elasticity is remarkably stable whether we group goods by port-month, commodity-month, port-year or commodity-year (estimates range from 0.082 to 0.096) or whether we aggregate observations by port over the year (0.085). Aggregating observations within a port over a year yields a similar estimate of 0.081. One should be cautious when comparing this monthly-port estimate to the estimate from the aggregated annual data as the samples are inherently different (indeed we have 1,753,642 groups of goods at the port-month level and 15,083 groups at the commodity-year level). Goods that are traded more frequently will be weighted more in our monthly sample.[17] In sum, with even more disaggregate data we still find sizeable pricing-to-market.

Our third and final approach to controlling for quality directly considers the quality bias in the data. In particular, the quality explanation for pricing-to-market essentially argues for an aggregation bias in which 10 digit categories are made up of different 11 digit goods that differ by quality and price, with rich countries buying relatively more of the relatively high quality and high price 11 digit goods than poor countries. Obviously we can not sign this bias from the 11 digit to 10 digit level; however, we can sign it at different levels, say from the 10-digit to 9-digit, or 10-digit to 8-digit and so on to examine the strength of the bias. (A formal justification of this approach, based

[16] For such a theory of the frequency of imports based on fixed ordering costs and inventory management see George Alessandria, Joseph Kaboski, and Virgiliu Midrigan (forthcoming).

[17] Aggregating goods by commodity, at the monthly level we find a coefficient of 0.148 while at the annual level it is 0.112. Eliminating the controls for frequency of shipment and mode of transport raise the elasticity to 0.184 at the annual level.

on a Lancasterian model where each Harmonized System digit represents a characteristic with a given price, and an explicit statement of the potential quality bias are developed in the unpublished web appendix.)

Table 4 reports the results of our aggregation exercise and finds that the pricing-to-market actually declines when we move from 10- to 9-digit, 10- to 7-digit, or 10- to 5-digit. Interpreting Table 4, in the 9-digit case, all categories that are identical up to the first nine digits are aggregated to construct 9-digit price data. Only 855 commodities are unique up to all ten digits and these are combined into 338 heterogenous 9-digit categories. As more digits are dropped, the categories become broader and more heterogeneous, more goods are combined into groups, and more observations can be included in the regressions. For example, at five digits artificial Christmas trees are simply artificial Christmas trees, while at seven digits these are subdivided into plastic and non-plastic artificial Christmas trees, one of which may have higher average prices and therefore be considered higher quality. (The unpublished web appendix contains a random selection of categories and how they are combined.)

In general, some but not all of the classification distinctions are driven by quality differences. For instance, olive oil is classified (at the sixth digit) into virgin (150910), refined (150990), and residue (151000), but knitted cotton sweaters are distinguished (at the ninth digit) by men's (611020101) and women's (611020102) with no obvious quality interpretation. To take our aggregation argument one step further, we identify a number of goods like olive oil, where there is clear quality variation, and redo the previous exercise. Specifically, we use identifiable quality differences in terms of freshness, size, purity, stage of processing, price and age (see details in the unpublished web appendix) to identify 1112 goods that can be paired with at least one more commodity. In total, we form 455 groupings. The bottom panel of Table 4 reports the results at the disaggregrate and aggregated levels treating each commodity-year observation separately (hence there are 3492 aggregated groups rather than 455). Using our disaggregated data we find a pricing-to-market coefficient of 0.151 and aggregating to the next level, the coefficient is 0.126. Thus, similar to our full sample of goods, for this targeted group of goods we also find that this form of aggregation bias lowers our pricing-to-market estimate.

One possible explanation for this result is that the bias goes the other way; although poor countries likely consume lower quality goods on average, conditional on importing from the United States, poor countries may import relatively higher quality goods. Indeed, one possibility is that since poor countries produce relatively low quality goods to begin with, they have better substitutes for low quality U.S. exports than other high income countries. That is, the U.S. may have a stronger comparative advantage in high-quality goods relative to poor countries than relative to other high-income countries.[18,19]

While we cannot conclude definitively that measured pricing-to-market in U.S. exports unit values is not due to quality differences, our various efforts to control for quality suggest that there is room for pricing-to-market from price discrimination. The model in the next section can be used to further distinguish the role of price discrimination vs. quality discrimination as an explanation for the pricing-to-market we find in the data.

D. Variation in PTM by Type of Good

We now examine which types of goods show the strongest pricing-to-market. In particular, our theory will emphasize search frictions as a source of market power and pricing-to-market, so we examine whether pricing behavior/unit values differ for goods where search costs may be highest.

[18] This comparative advantage interpretation of our result for U.S. exports is consistent with Schott (2004), who finds that U.S. imports coming from lower income countries tended to be less expensive. Since he studies goods exported from multipe sources into a single market (the U.S.), he attributes this price variation to quality variation rather than PTM. If the U.S. has a stronger comparative advantage in high quality goods production relative to poor countries, then one would expect U.S. imports from poor countries to be lower quality goods, while its exports are higher quality.

[19] Another possibility is that per-unit trade costs are relatively higher for poor countries, and hence they import relatively higher quality goods than in rich countries (see David Hummels and Alexandre Skiba 2004)

We measure the importance of search for a good in three ways: (1) distinguishing goods by final purchasers; (2) method of sale; and (3) distinguishing goods by the importance of repeated transaction/long term relationships. For all three, we find that higher search costs are consistent with more pricing-to-market.

End Use

We first distinguish goods by end use. The opportunity cost of search might also matter to firms in their decisions to search (see James E. Rauch, 1999, 2001, Alessandra Cassella and James E. Rauch, 2003, and James E. Rauch and Joel Watson, 2003), but we model consumer search, and the story applies most naturally to consumers.[20]

We do indeed find the strongest results for consumer goods. Table 5 reports our coefficients by end-use category (1-digit codes). Consumer goods have the highest estimated income elasticity at 0.218. The other four categories (we exclude re-exports and "other") are all positive but lower than consumer goods, averaging just 0.130.

Method of Sale

A second way to measure the importance of search costs is to distinguish goods by their method of sale. To test the role of search frictions on trade flows, Rauch (1999) classifies goods into three classes: 1) Organized exchange, 2) Reference priced and 3) Differentiated goods. The Rauch classification is based strictly on the dominant method of sale for goods within a product category. In principle, price data is most readily available for goods traded on a central exchange and least available for differentiated goods which require search between suppliers and purchasers.[21] The lower panel of Table 5 presents the estimates of pricing-to-market by a good's mode of sale.

[20] In all likelihood, few of these exports are sold directly to consumers. However, the end user is more important than the purchaser, since the end user will determine the elasticity of wholesaler demand (see the unpublished appendix).

[21] Tibor Besedes and Thomas J. Prusa (2006) also use the Rauch classification to infer differences in search costs across goods.

We find the highest price-wage elasticity for differentiated goods of 0.163 and the lowest for goods sold on organized exchanges of 0.074 percent.

While the Rauch classification targets search costs directly, and so this evidence may appear most direct, there are several caveats. First, the Rauch classification largely picks up differences in the end-use classification system. Controlling for end-use classification, there are much smaller differences in pricing-to-market by Rauch classification (Table 5, lower panel). For the most part, "organized exchange" goods picks up agricultural and industrial supplies while "differentiated" goods are more likely to be consumer or capital goods. Second, although the classification is based on method of sale, not the similarity of goods within a product category, more standardized goods are easier to be traded on organized exchanges or reference priced. Hence, these classifications may also relate to the heterogeneity (of which quality is only one dimension on which goods differ) of each good. Indeed, if one expects the law of one price to hold on organized exchanges then the coefficient here may be troubling, and one might expect that even our organized exchange goods are highly heterogeneous. However, the third caveat is that the classification is very coarse, classifying only 1190 categories at the 4-digit SITC level rather than the 10,304 categories at the 10-digit HS level we use. That is, each category may include HS-10 goods sold in all three manners, but an organized exchange category should contain relatively more goods sold on organized exchanges, for example. Finally, with regards to organized exchanges, the fourth caveat is that, even within Rauch's "conservative" classification, a large share of these products are not actually transacted centrally on organized exchanges. Futures are available for these goods, but transactions are purely financial based on reference prices or price indices over differentiated goods. Indeed, futures are structured in this way because the commodities are not centrally traded and not standardized.

To address these last three issues, we constructed a finer classification based on actual contract details that directly maps centrally transacted goods into 10-digit HS goods. We used all 44 goods transacted on U.S. commodity exchanges over the relevant period (see unpublished web appendix for details). Using this narrow classification, we have 4253 observations. We get a coefficient of 0.023 (with robust standard error of 0.016) on the log wage, which is substantially lower than our coefficient for the full sample of 0.165, and it is not statistically significant at even the 10 percent level. Hence, our data measure little to no pricing-to-market for centrally transacted goods. We interpret this as resulting from a lack of market power because of price availability, but given our quality caveat, it could also be a result of no heterogeneity among these goods.

Repeated Transactions

Search is likely to be less important when goods are sold as part of repeated transactions or long term relationships. However, controlling for repeated transactions and long term relationships is difficult.[22]

One way is to distinguish, at the good level, which goods are more or less likely to be sold in repeated transactions. Some goods are commonly made "to order", according to detailed customer specifications, and these may be part of a long term relationship that involves repeated transactions.[23] Other goods are commonly produced "to stock", i.e., without a customer in mind, and placed directly into inventory. David A. Belsley (1969) and Mark Bils and James Kahn (2000) note that the latter are likely to have higher inventory/sales ratios, and they classify two-digit SIC sectors in this way. Another way of making the same distinction is to look at the share of final goods in total inventory (measured using 1997 Census of Manufactures data for 5- and 6-digits NAICS industries). If a product is to order then presumably the producer has little incentive to hold goods in final inventory when there is a customer already in place (indeed the Belsley/Bils/Kahn to stock

[22] An alternate possibility is that export pricing is due to mismeasurement from transfer pricing considerations by related parties. In our working paper (George Alessandria and Joseph Kaboski 2007) we find that controls for transfer pricing considerations do not change our estimates.

[23] We thank Steve Davis for suggesting this distinction to us.

industries hold about 1/3 more final goods than to order industries). We classify the commodities with a high share of final inventories (the top quartile) as "to stock" industries.

Table 6 shows that, using all three measures of goods, pricing-to-market coefficients are higher (0.185 to 0.207) for "to stock goods," where repeated transactions are less likely. This is supportive of a role for search, but we also find substantial pricing-to-market for "to order" goods where repeated transactions are common of 0.146 to 0.160. This may indicate that things other than search drive pricing-to-market or that, conversely, search may matter even for repeated transactions. Alternatively, it may indicate that our inventory distinctions are a weak metric for repeated transactions, perhaps because the distinction for domestic transactions carries over only weakly to international transactions.

Another way to account for the importance of repeated transactions is to measure and control for them directly at the customer level. We do this using the monthly U.S. export data. Specifically, we take one year of data, 1999, and then construct commodity-country unit values for each quarter within that year. We then estimate pricing-to-market using unit values in the fourth quarter with controls for the number of quarters with prior transactions,

$$(3) \quad \ln p_{ij} = \beta \ln w_i + \alpha_1 I_{1ij} + \alpha_2 I_{2ij} + \alpha_3 I_{3ij} + \varepsilon_{ij},$$

where I_{1ij}, I_{2ij}, I_{3ij}, are separate dummies for one, two or three quarters of transactions.

Table 7 reports that with these controls for repeated transactions our pricing-to-market coefficient falls from 0.215 to 0.194 and that the discount for more frequent buyers is significant, sizeable, and increasing in the number of previous transactions. A destination that imported from the U.S. in one previous quarters gets a 4.9 percent discount while a destination that purchases in all three quarters gets a 34.5 percent discount. That our pricing-to-market estimate is not substantially altered by the presence of large relationship specific discounts suggests that the structure of U.S. trade relations does not differ dramatically by income. Additionally, the large discounts on repeat transactions provides some further evidence against the quality variation interpretation of the price variation as there is no obvious reason to suspect buyers of low quality goods are more likely to be in long-term relationships.

E. Importance of Pricing-to-Market

Our first aim was to estimate the importance of pricing-to-market for deviations from absolute PPP. To do so, we modify Engel's, now standard, decomposition of fluctuations in real exchange over time to take into account differences in price levels across countries by income,[24]

$$\frac{p_i - \bar{p}}{y_i - \bar{y}} = \frac{p_i^T - \bar{p}^T}{y_i - \bar{y}} + (1-\alpha)\frac{\left(p_i^{NT} - p_i^T\right) - \left(\bar{p}^{NT} - \bar{p}^T\right)}{y_i - \bar{y}},$$

$$\Rightarrow \quad \varepsilon_{PPP} = \varepsilon_{LOP} + (1-\alpha)\,\varepsilon_{N/T}, \tag{4}$$

where ε_{PPP} is the elasticity of the overall price level (with respect to income per capita), ε_{LOP} is the elasticity of deviations from the LOP (in tradables), $\varepsilon_{N/T}$ is the elasticity of the relative price of non-tradables, and α is the share of tradables. This decomposition shows that the aggregate price level and income relationship we seek to explain, ε_{PPP}, depends one-for-one on the deviations from the LOP and income relationship, ε_{LOP}, and only partly on how the relative price of non-tradables to tradables varies with income across countries, $\varepsilon_{N/T}$.

In the data of Figure 1, $\varepsilon_{PPP} = 0.48$. The traditional HBS story assumes $\varepsilon_{LOP} = 0$ so that the 100 percent of these deviations from PPP come from the relative price of non-tradables to tradables. But the data tell us there are sizeable deviations from the LOP. The PWT data on the

[24] This assumes that the log price index is approximated by a geometric average $p = \alpha p^T + (1-\alpha)\,p^{NT}$ and that all countries have the same basket.

price of tradables indicate that $\varepsilon_{LOP} = 0.31$, so that deviations from the LOP account for about 65 percent of the aggregate price-income relationship. The more modest estimates of pricing-to-market from the export data of $\varepsilon_{LOP} = 0.23$ would still account for nearly half of the PPP income relationship. These findings are consistent with Engel's finding that about 45 percent of U.S. long-run real exchange rate fluctuations are due to movements in the relative price of traded goods.[25] We interpret the difference between the two values for ε_{LOP} (0.31-0.23=0.08) as measuring the 17 percent (0.08/0.48) contribution of local non-traded distribution costs to differences in price levels across countries.

II Search as a Theory of Pricing-to-Market

This section develops a search-driven theory of pricing-to-market, in which firms charge high prices on average in countries where wages, and hence the opportunity costs of search, are high. Consumers in high-wage countries are less willing to spend time searching for low prices. The theory produces a positive relationship between prices, wages, and income. We first discuss evidence in support of such a theory and then present a formal model.

A. Corroborating Support for Search

There are a number of reasons to favor the search-based story over a direct preference story in which consumers become less price sensitive with income. First, there is substantial evidence within countries that prices are dispersed and that consumers alter their shopping behavior to take advantage of this dispersion of prices. For instance, Aguiar and Hurst (2007), using scanner data on consumer expenditures and diary data on time-use, find that a doubling of shopping time lowers the average purchase price by 7 to 10 percent.[26] Second, within countries there is evidence that

[25] Engel (1999) attributes approximately 95 percent of *short-run* real exchange rate fluctuations to movements in the relative price of traded goods.

[26] Shopping with uncertainty, either due to the availability of the goods or time of the shopping trip, is isomorphic to a model with no uncertainty over search time but uncertainty over prices.

shopping effort, measured as time spent shopping per dollar spent, is decreasing in the wage of shoppers. Third, cross country evidence on time-use suggests that low-income countries have a comparative advantage in producing non-traded search services. Fourth, empirically, the evidence from our U.S. export pricing-to-market estimates are consistent with the search story. A final methodological reason is that the search story offers a true explanation of this relationship, rather than just assuming it through preferences.

Recall that several pieces of evidence from the U.S. export data are consistent with the search explanation. First, it is the wage level rather than income per capita that drives the pricing relationship, when both explanatory variables are included. In the search story, the elasticity is driven by the opportunity cost of time (i.e., the wage) rather than non-labor income or differences in income per capita arising from demographic differences. Admittedly, it is possible that the significance of wages could be driven by other factors, such as measurement quality or coverage of the wage data (which is strictly the manufacturing wage). Second, goods for which search costs are highest display the most pricing to market.[27]

Time-use studies provide evidence that time spent shopping is related to income in a way consistent with the search model. Many studies have examined the relationship between the opportunity cost of time and shopping behavior for consumers within a given economy, i.e., facing a given distribution of prices. For example, McKenzie and Schargrodsky (2005) study the behavior of Argentinian shoppers and find a strong relationship in the cross-section of consumers between consumer search and income. After controlling for quantity purchased, they find that consumers in

[27] One further piece of suggestive evidence from these data is that high wage countries do not simply pay higher average prices; there is also evidence that they face higher dispersion in prices. Specifically, unit value dispersion (calculated as the standard deviation of log unit values across all 10-digit disaggregated HS codes within a higher level of aggregation, i.e., 9-digit, 8-digit, and 7-digit levels for a given country) in the annual U.S. export data is higher for high income countries. The elasticities of log unit value dispersion with respect to log real wages range from 0.041 (for 9-digit) to 0.047 (for 7-digit), and these results are marginally significant. Aggregation greatly reduces sample sizes – recall Table 4 – and so t-statistics range from 1.5 to 1.8. When using exchange rate-converted wages rather than PPP-converted wages, the estimates are somewhat larger and significant even at the 5 percent level.

the 10th percentile of the income distribution spend 30 percent more time shopping than consumers in the 90th percentile. Low-income consumers also visit a greater variety of stores. McKenzie and Schargrodsky also show that the 2002 Argentine economic crisis, which lowered the wages of workers, led to increases in both these measures of consumer search. Still, financial crises presumably affect both income/wealth and the distribution of prices, in addition to the opportunity cost of time. Aguiar and Hurst (2007) have cleaner evidence of the effect of opportunity cost of time. They document an increase in shopping time per purchase experienced upon retirement, which affects the opportunity cost of time, but should not affect the lifetime budget constraint nor the distribution of prices. Both of these studies also find that search effort is negatively related to purchase price. This evidence of dispersion in prices even within countries provides further support for our search story over one based on tastes.

The cross-country time-use data are also consistent with our theory. Since the distribution of prices is not the same over the cross-section of countries, the search story does not (necessarily) imply that consumers in poor countries shop more per unit purchased than consumers in rich countries. (Indeed, in the model, the response of firm pricing behavior exactly cancels out the increased willingness to search in poor countries, and search time per unit is constant across countries. All differences in search effort work through the reservation price consumers are willing to accept.) However, our story hinges on the cost of shopping rising with country income, which requires productivity in the production technology to rise faster than the productivity in the shopping technology. Since income (and purchases) rises faster than shopping productivity, a crucial implication of the theory is that people in rich countries spend more total hours shopping per hour of work than people in poor countries.

In Table 8 we report the relationship between time-use and income per capita from two separate cross-country time-use datasets. The first line reports the results from the recently completed European Harmonized Time-Use Survey (EHTUS). The second line reports the results from the Multinational Time-Use Survey (MTUS). In general, cross-country time-use comparisons are difficult owing to definitional and sampling differences. The EHTUS was designed with these comparability issues in mind, while the MTUS is a collection of mostly individual country surveys that have been recoded to be more comparable ex-post. Despite these differences, we find that both surveys generate a similar relationship between the ratio of shopping to work time and income per capita. From the EHTUS we find a 10 percent increase in income per capita generates a 3.4 percent increase in the ratio of shopping to work time (3.2 percent in the MTUS).

B. Model

To generate pricing-to-market, exporters must have some market power when selling into a particular destination and there must be some barriers to international arbitrage. In our framework, we model many exporters directly selling an identical product directly to consumers. Even though each exporter sells the same product, the search frictions give each exporter some market power. The extent of the market power, and thus the properties of the demand curve, is given by the structure of search. Search also makes international arbitrage costly. For simplicity, we omit a separate retail and distribution sector, but, we have shown that our results are robust to the inclusion of a separate retail and distribution sector, where consumers search among retailers that purchase a differentiated good from producers because producers (and exporters) take into account downstream consumer search behavior when setting prices to retailers. (See the unpublished web appendix.)

Environment

There are three imperfectly substitutable goods $i = \{1, 2, 3\}$ and two countries denoted $j = \{1, 2\}$. Goods 1 and 2 are tradables, with good 1 produced exclusively in country 1 and good 2 produced exclusively in country 2. Both countries can produce good 3, but it is not tradable.

Including non-tradables along with tradables allows us to incorporate, and distinguish between, the traditional HBS effect and pricing-to-market.

In each country, there are many stores, each specialized in the sale of a single good. For simplicity, we assume that the measure of each type of store in each country is predetermined and the same. Households do not know the price charged at any store and must physically visit a store to discover its price. Because search takes time and is imprecise, stores have some monopoly power over consumers and thus may charge different prices for the same good. We assume stores are owned and operated by the firm producing output, but require no additional inputs. We abstract from wholesale, retail, and international trade costs since we found them to be only half as important as pricing-to-market for tradable prices.

Households send out shoppers to search for the lowest price quotes and purchase goods. Each shopper can buy at most one unit of the good. Shopping therefore takes time away from work and is imperfect in the sense that consumers do not simultaneously receive price quotes from all the stores in the market. We model search as noisy, as in Kenneth Burdett and Kenneth L. Judd (1983), so that a fraction q of shoppers receive a single price quote while the remaining shoppers $(1-q)$ receive two price quotes. The probability that a shopper receives a single price quote is random and equals q. After receiving either one or two price quotes, the shopper must decide whether to purchase a single good at the lowest price quote received or return home empty-handed.

Although without searching agents do not know the price charged at a specific store, they do know the distribution of prices in the economy. A shopper from country j looking for good i receives (domestic) price quotes for good i from the known distribution $G_{ij}(.)$. Since the shopper can buy at most one unit of the good, only the lowest price quote received by a shopper is relevant to the shopper's purchase decision. The distribution of lowest price quotes is then

$$H_{ij}(p) = qG_{ij}(p) + (1-q)\left[1-(1-G_{ij}(p))^2\right].$$

From the firm's perspective, noisy search makes the consumers heterogeneous in that some shoppers will have only one price quote, while others will have multiple price quotes. Consumers with multiple price quotes will differ in their second price quote. Since firms cannot distinguish between these different customers, the price they charge will influence both the profit per sale and the share of shoppers with multiple price quotes that purchase from them.

Consumer's Problem

The consumer's problem is similar to that in George Alessandria (2009). In each country, there are many identical families. Lowercase variables denote individual decision rules and uppercase variables denote aggregate decision rules. Each family is composed of a large number of agents, normalized to a continuum of measure one. The problem of a family is to divide its agents between working and shopping and to give shoppers instructions on which prices to accept. In country j the number of agents n_{ij} shopping for good i and the number of agents l_j working satisfy the time constraint:

$$(5) \quad \sum_i n_{ij} + l_j = 1.$$

It is optimal to send each agent shopping for good i with a reservation price rule to purchase only if the lowest price quote is below some reservation level, r_{ij}. Consumption of good i by country j consumers depends on both the reservation price and the measure of shoppers. With many shoppers for each good there is no uncertainty in consumption, which equals:

$$(6) \quad c_{ij} = n_{ij} H_{ij}(r_{ij}).$$

Given the reservation price, the average purchase price is evaluated from the truncated distribution of lowest prices:

$$(7) \quad p_{ij}(r_{ij}) = \frac{\int_0^{r_{ij}} p dH_{ij}(p)}{H_{ij}(r_{ij})},$$

which is clearly increasing in reservation price.

The representative home family chooses reservation prices and shoppers for each good to solve the following problem:

$$\begin{aligned} U^j &= \max_{\{r_{ij}, c_{ij}\}} U(c_{1j}, c_{2j}, c_{3j}), \\ \text{subject to} &: \begin{cases} \sum_i p_{ij}(r_{ij}) c_{ij} = w_j l_j + \Pi_j, \\ \text{equations (5), (6), (7),} \end{cases} \end{aligned}$$

where U^j is the utility function in country j and Π_j is the profits earned by country j firms.

If an interior solution exists the first-order conditions satisfy:

$$\begin{aligned} (8) \qquad r_{ij} &= \frac{w_j}{H(r_{ij})} + p_{ij}(r_{ij}), \qquad i = 1, 2, 3, \\ (9) \qquad \frac{U_1^j}{U_i^j} &= \frac{r_{1j}}{r_{ij}}, \qquad i = 2, 3, \end{aligned}$$

where U_i^j is the marginal utility of good i.

Equation (8) is an arbitrage condition that implies, at the margin, the family is indifferent between (1) increasing consumption by purchasing at the reservation price, or (2) sending out additional shoppers, whose opportunity cost of search is measured in terms of the forgone wage, and purchasing at the average price. With a reservation price of r_{ij}, the family expects to send out $1/H\left(r_{ij}\right)$ shoppers to purchase a single unit. Since the reservation price is linked to the true cost of the good, this is the cost that matters at the margin; therefore, the family chooses consumption so that the marginal rate of substitution between any two goods equals the ratio of their reservation prices as in equation (9).

We focus on the difference in prices across countries with different incomes and therefore only consider a representative agent in each country. However, it is straightforward to extend the model we present to permit heterogeneity in wages. In this case, we see from equation (8) that within countries, consumers with relatively high wages will have high reservation prices, and search less intensively, than consumers with relatively low wages, consistent with the within-country evidence.

Firm's Problem

There are many firms in each country that specialize in the production of either the country's tradable or non-tradable good. Firms within a country are *ex ante* identical. Labor is the only input into production, and one unit of labor in country j produces a_j^T units of the tradable good (good j) and a_j^{NT} units of the non-tradable good (good 3). To focus on international price discrimination, firms can costlessly sell their goods in either country through the pre-established outlets.

To fix ideas, consider the problem of a representative firm in country 1 selling the tradable good (good 1) in country j. A similar problem exists for non-tradable and country 2 firms. Even though many firms produce the same good, the search frictions give each firm some monopoly power and lead firms to behave as monopolistic competitors. Each firm takes as given the distribution of prices charged by other firms selling the same good, G_{1j}, the number of price quotes that it delivers, the reservation price of consumers, R_{1j}, and the unit cost of production, w_1/a_1^T. Given the constant returns to scale production, the amount of sales does not influence a firm's unit cost. Thus, the firm's problem becomes one of maximizing profits per customer that receives a price quote. The representative firm from country 1 selling in country j solves:

$$\pi_{1j} = \max_{p} \left(p - \frac{w_1}{a_1^T} \right) Q_{1j}(p),$$

where $Q_{1j}(p)$ is the probability that a firm makes a sale when charging a price p and equals:

$$Q_{1j}(p) = \begin{cases} \frac{q}{2-q} + \frac{2(1-q)}{2-q} \left[1 - G_{1j}(p)\right] & \text{for } p \leq R_{1j}, \\ 0 & \text{otherwise.} \end{cases}$$

As long as the firm's price is below the reservation price, the firm will sell to all customers with one price quote. By increasing its price, the firm increases its revenue per sale but decreases the likelihood of a sale, since it increases the probability that those customers with two price quotes have a second price quote that is lower than the firm's price.

Burdett and Judd (1983) show that given a reservation price, R_{ij}, and cost of production, w_i/a_i^T, a unique distribution of prices exists, $G_{ij}(p)$, where

$$G_{ij}(p) = \begin{cases} 0 & p < \underline{P}_{ij} \\ 1 - \frac{q}{2(1-q)} \frac{R_{ij} - p}{p - w_i/a_i^T} & p \in \left[\underline{P}_{ij}, R_{ij}\right] \\ 1 & p > R_{ij} \end{cases} \quad \text{and} \quad \underline{P}_{ij} = \frac{2(1-q)\, w_i/a_i^T + qR_{ij}}{2-q}$$

Any price on the support of the distribution yields firms the same profits, and firms will randomize. Firms with relatively high prices primarily sell to those consumers with a single price quote, while those with relatively low prices attract more of those shoppers with multiple price quotes.

Equilibrium

The total demand for labor by firms producing tradables and non-tradables in country j are L_j^T and L_j^{NT}, respectively. The labor market clearing condition is:

$$L_j^T + L_j^{NT} = \frac{N_{j1} + N_{j2}}{a_j^T} + \frac{N_{3j}}{a_j^{NT}} = L_j$$

A symmetric equilibrium is then a distribution of prices, G_{ij}, and wages, w_j; consumer decision rules $\{l_j, n_{ij}, r_{ij}\}$ and aggregate decision rules $\{L_j, N_{ij}, R_{ij}\}$ in each country $j = \{1, 2\}$ for each good $i = \{1, 2, 3\}$ such that: (1) Given prices, wages, and profits, consumer's decision rules solve the household's problem in each country; (2) Given prices and wages, each firm chooses a price to solve each firm's problem; (3) Goods and labor markets clear; and (4) Individual and aggregate decisions are consistent so that all households from the same country behave identically.

Alessandria (2009) shows that the highest price in the market equals the reservation price. This upper bound on prices is an equilibrium because the highest-priced firms have no incentive to charge a price above the reservation price, as they would lose all sales. As no shopper returns empty-handed, the marginal cost of each good in each country is the average price paid for it plus the opportunity cost of the shopper. This equals the reservation price:

$$r_{ij} = w_j + p_{ij}(r_{ij}).$$

We focus only on the average transacted price (which equals the unit value), since this most closely corresponds to the measure used by the national statistical agencies and in our empirical work. Substituting the equilibrium reservation price into the distribution of prices, we can solve for the average price for tradables of good i (from country i) and non-tradables sold in country j as:

$$(10)\quad p_{ij} = \frac{w_i}{a_i^T} + \frac{q w_j}{1-q}.$$

$$(11)\quad p_{3j} = \frac{w_j}{a_j^{NT}} + \frac{q w_j}{1-q}.$$

The average price for good i paid by a consumer in country j is equal to a markup over the marginal cost of the firm from country i. The markup depends on both the information structure of search (summarized by q) and the time cost of search w_j. Holding q constant, agents in a country with a low wage will, on average, pay a lower price than agents in a country with a relatively high wage. Consequently, the model predicts a strong relationship between prices and local wages.[28]

III Results

This section evaluates the model's quantitative properties. We first show that the model generates the relation between prices, shopping time, and wages within countries documented by Aguiar and Hurst (2007). We then show that the model generates large deviations from the LOP across countries even when productivity differences across countries are the same in tradables and non-tradables. Moreover, we find that our model closely matches the observed relationship between wages and tradable prices. We then examine the importance of pricing-to-market relative to the traditional HBS effect arising from productivity differences biased toward tradables.

[28] Equation (10) clarifies the difference between our model and HBS. In both models, tradables may sell for different prices across countries. In HBS, the price of tradables may differ internationally when there is a non-traded input, such as wholesale or retail distribution, to get the good to the final consumer. In our model of pricing-to-market, the search cost is similar to the non-traded retail or distribution costs in HBS. Unlike in HBS, this search cost is borne by the consumer and through the search frictions it is incorporated into the price charged at the border.

A. Calibration

Preferences are consistent with the standard textbook presentation of the HBS model.[29] Agents in each country have the following symmetric utility function:

$$U^j = u\left(c_{1j}, c_{2j}, c_{3j}\right) = \left(c_{1j}^{\rho} + c_{2j}^{\rho}\right)^{\frac{\alpha}{\rho}} c_{3j}^{1-\alpha}.$$

Preferences over tradables and non-tradables are Cobb-Douglas.[30] Home and foreign tradables are often assumed to be perfect substitutes. We depart slightly from this case and set $\rho = 0.99$.[31]

The size of the tradable sector is set to match the median trade share of GDP of those OECD countries for which re-exports are not large[32] and for which we have manufacturing wage data from the BLS. The median country[33] imports and exports approximately one-third of GDP in 2000 and so we set $\alpha = 2/3$ and non-tradables account for one-third of output.[34] The openness of a country affects the weight we put on the HBS channel but does not substantially change the amount of pricing-to-market. We report sensitivity to the trade share.

The production side of the economy is calibrated as a symmetric two-country model to match certain features of the U.S. economy. The production parameters are the search, q, and market goods, a^T, productivities. For our baseline case, we assume that tradable and non-tradable technologies are identical,[35] so that $a_j^T = a_j^{NT} = \bar{a}$. Since productivity in market shopping is 1 (each shopper can purchase one unit), $\bar{a}$ captures the relative productivity of market production to

[29] See Maurice Obstfeld and Kenneth Rogoff (1996).

[30] The assumption of a unitary elasticity of substition between tradables and non-tradables is consistent with the estimate of 1.24 by Jonathan Ostry and Carmen Reinhart (1991) for a group of developing countries and Enrique Mendoza's (1995) estimate of 0.74 for a group of industrialized countries.

[31] Since our focus is on the long-run differences in price levels, our calibration of ρ differs substantially from models focused on short-run fluctuations.

[32] This requires dropping the Netherlands, Belgium, and Ireland.

[33] For comparison, the median country in the Penn World Tables imported approximately 38 percent of GDP and exported 42 percent of GDP in 2000.

[34] Alan Stockman and Linda Tesar (1995) use data on a cross-section of OECD countries from 1970 to 1985 and find the tradable sector is nearly 50 percent of output.

[35] The model is calibrated to the typical good. A more general model would allow for goods to vary in both the noisy search parameter and the time it takes per purchase, while holding these parameters constant across countries.

shopping. In equilibrium, since all produced goods are purchased: $\bar{a} = \frac{N}{L}$, the ratio of shopping time to market labor. The American Time-Use Study (2003) reports that the average American spends about 4 times as much time working as purchasing goods and services, so $\bar{a} = 1/4$. The labor's share parameter θ is set to 60 percent of total income (Thomas F. Cooley and Edward C. Prescott 1996), and this pins down $q = 0.727$. In all of our experiments, we hold this noisy search parameter constant across countries but allow tradable and non-tradable productivity to vary.[36]

We use the model to construct a distribution of prices and income, which we then compare to the data. We do this by solving our two-country model repeatedly. In each case, one country is the U.S., and the second country is a PWT benchmark country. Productivity in the second country is chosen to match income per capita relative to the U.S. In this way, we match the world income distribution and have synthetic price data for 115 artificial economies.

With two symmetric countries, our calibration implies an average markup over marginal cost of 66 percent. However, because exporters reduce their markup to low-income countries, in the asymmetric version of the model, the average markup of firms from the richest country is only 60 percent, and 40 percent on average across the 115 countries. Markups are notoriously difficult to measure, yet this level of monopoly power is consistent with those found in structural IO studies of the ready-to-eat cereals market (Avi Nevo 2001) and U.S. automobile market (Steven Berry, James Levinsohn and Ariel Pakes, 1995, and Pinelopi Goldberg, 1995)[37]. Moreover, our pricing-to-market evidence finds on average there is a 48 percent price difference for the same goods between the richest and poorest countries. Such price variation is only possible if markups are of this size.

[36] The choice of q will determine the markup and will influence the slope of the price-income relation across countries. However, q is calibrated independently of its implications for the slope.

[37] A typical markup in macro papers is about 30 percent (see Michael Dotsey, Robert King and Alexander Wolman, 1999, and Andrew Atkeson and Ariel Burstein, 2008). Estimates of markups from the structural IO literature tends to find markups that are much higher. For comparison, Berry, Levinsohn, and Pakes (1995) estimate that markups range between 31 and 60 percent. Goldberg (1995) which specifically focuses on international trade in cars estimates markups for automobiles are on average 61 percent. Nevo (2001) who studies the ready to eat cereal industry finds even larger markups, ranging from 50 percent to 110 percent.

We first consider the results when the productivity gap between countries is the same in both sectors. This is our *Balanced Productivity* gap case. In our benchmark, the *Biased Productivity* gap case, we assume that productivity difference in the non-tradable sector is smaller than in the tradable sector. While it is commonly asserted that the productivity gap in tradables is relatively large compared to non-tradables, there is little direct cross-country evidence of this gap. Studies that do measure this gap across countries assume that the LOP holds for traded goods and use relative prices to infer productivity differences.[38]

Rather than use our theory to construct relative productivity differences, we consider the evidence on the relationship between income and productivity in tradables and non-tradables, respectively, in the U.S. time series. Dale W. Jorgenson and Kevin J. Stiroh (2000) estimate labor productivity growth by industry for the U.S. from 1958 to 1996. We split these industries into tradable and non-tradable sectors and then construct a measure of the productivity gap using each industry's share of sectoral value-added. These weighted averages, along with simple averages, of TFP and labor productivity growth rates[39] are reported in Table 9. We find that non-tradable labor productivity has grown about two-thirds as fast as tradable labor productivity. In the biased productivity case, we take the time-series evidence from the U.S. on the productivity gap and examine the implications of such a gap for the world distribution of income and prices. We also test the sensitivity to the relative size of this gap. For reference, we also present results from the standard HBS model with no pricing-to-market.[40] Table 10 includes the parameters for the various models.

Prices and income are measured consistently with the empirical data and statistics computed

[38] For instance, Hsieh and Klenow (2007) use the relative price of consumption to investment to infer that productivity in the investment sector increases with output (in the cross-section) at a rate 2.6 times that of the consumption sector. Similarly, using data on relative price levels, Berthold Herrendorf and Akos Valentinyi (2006) find that the productivity difference in tradables must be nearly 12 times larger than those in non-tradables.

[39] We follow Mathew Canzoneri et al. (1999) and focus on labor productivity. What really matters for the HBS effect is the change in the marginal product of labor across sectors. For a broad range of production functions this is proportional to the change in average labor productivity. In contrast, measuring TFP growth depends on the assumed structure of the production function and requires measures of capital stocks and materials usage.

[40] Our model converges to the HBS model as $q \to 0$ and a becomes large.

previously. Deviations from the LOP are measured as the log average price of U.S. exports to destination j, or $\ln LOP_j = \ln\left(P_{U.S.,j}/P_{U.S.,U.S.}\right)$. To measure income, we follow the convention of the Penn World Tables and compute nominal GDP, Y_j, as the sum of expenditures of domestic production. The aggregate price level, P_j, is measured using the welfare-based price index.[41] Real income, y_j, is nominal GDP deflated by the price index P_j (i.e., $y_j = Y_j/P_j$).

With these measures of real income and prices, we estimate statistics that correspond to our empirical results.[42] All results are presented in Table 11. The table's top panel presents our elasticity estimates from various versions of the model. The bottom panel decomposes each model's price-income relation into its main components. The column titled *Data* summarizes our estimates of the price-income relationship from the PWT tables (ε_{PPP} and $\varepsilon_{N/T}$) and the U.S. export data (ε_{LOP} and ε_w), plus evidence from the time-use surveys $\left(\varepsilon_{shop/work}\right)$.

B. Prices, shopping, and wages within markets

Before considering the cross-country implications of our model, we show the model is consistent, both qualitatively and quantitatively, with some key features of prices paid, shopping time, and opportunity cost of time within countries documented for U.S. consumers by Aguiar and Hurst (2007).[43] Within countries our model makes two key predictions: first, increasing search lowers prices paid; and second, increasing the opportunity cost of search lowers search intensity.

With respect to the tradeoff between search and prices paid, Aguiar and Hurst find that a doubling of shopping time reduces the average price paid by between 7 and 10 percent. While all agents are identical in our model, we can study the behavior of an individual searches more than the representative agent by adjusting the wage, and therefore the reservation price, of a single individual

[41] The price index takes into account only the transaction price of the goods, not the search costs that are borne. Deriving price indices that include the costs of search do not noticeably change the quantitative results. Also, measuring output at world prices generates similar results.

[42] We compute a single price for each country and then estimate our statistics from the synthetic sample of countries running the same regressions we ran in the empirical section.

[43] We thank an anonymous referee for suggesting rigorously analysizing the within country predictions of our model.

but keeping the distribution of prices fixed. In our benchmark calibration for the U.S., we find that a single consumer that shops twice as often as the representative agent, ends up paying 5.5 percent less on average than the representative consumer.[44]

With respect to the relation between the opportunity cost of time and search effort, Aguiar and Hurst present convincing evidence that high opportunity cost individuals search less intensively. In our model, the opportunity cost of search is linked to the market wage and so we can ask: how does search time vary with the opportunity cost of time within countries? We find that an agent with half the median wage searches nearly 20.4 percent more and pays 2.4 percent less than someone with the median wage. Aguiar and Hurst lack data on wages and instead proxy for the opportunity cost of time with household characteristics. They find search effort is higher for retired people, single-earner households, smaller households, and poorer households. Without a precise measure of the opportunity cost of time, it is difficult to precisely estimate how search effort varies with opportunity cost of time. However, when looking at different groups by income, an imprecise measure of the opportunity cost of search, they find high income consumers pay 2.1 percent more and search 16 percent less. Thus, there is strong support for the two key predictions of the model.

Aguiar and Hurst also directly measure the price-income relation within countries. They find evidence that prices paid increase with income: consumers with income below $30,000 pay 2.1 percent less than those with income above $70,000. Converting these income brackets to averages, the data suggest a price income relation of about 1.2 percent. In our model, considering a consumer with half the wage, we find a 3.1 percent price income relation. We make this comparison with caution, since we do not have a rich model of the relationship between opportunity cost of time and income in a population cross-section.[45]

[44] Generating a larger reduction in prices from search, at the midpoint (8.5 percent) of the Aguiar and Hurst range, requires increasing the opportunity cost of shopping, which we can do by lowering the work to shopping time ratio in half and adjusting q to maintain the same labor share (q=0.57).

[45] Aguiar and Hurst results suggest that demographics are important, for example. Another issue is that search

Given the small effect of income/wages on prices paid within countries, one might suspect that our model could only explain a small fraction of the price differences across countries. However, this logic ignores the general equilibrium effect of firms adjusting prices by market. Specifically, firms facing a single low wage consumer in a market will not alter their prices in response to this consumer's reservation price strategy, but will adjust their prices when facing a market full of low wage consumers. Consider again a consumer with half the median wage in the U.S.. This consumer has a 7.4 percent lower reservation price and comes home empty-handed 20.4 percent of the time. Facing a market full of low wage consumers, however, firms lower their prices in order to attract more consumers. Consequently, in general equilibrium a representative consumer in a country with half the median U.S. wage will pay 22.3 percent less than U.S. consumers.[46]

C. Balanced Technology Gap

In the balanced technology case, the productivity gap across countries is the same in both sectors, $a_1^T/a_j^T = a_1^{NT}/a_j^{NT}$. These cross-country productivity differences generate differences in wages, income and prices. In equilibrium, this generates higher prices for all goods, tradable and non-tradable, in higher wage/income countries. With tradables accounting for two-thirds of expenditures, the model generates a quantitatively significant amount of pricing-to-market. The elasticity of deviations from the LOP with respect to wages is 0.115, which is over 50 percent (0.115/0.209) of what we observed in the data. Moreover, the elasticities of deviations from the LOP (0.114) and violations of absolute PPP (0.107) are also substantial in the model. Thus, our model can account for about half (0.114/0.235) of the deviations from the LOP and almost 25 percent (0.107/0.480) of the violations of PPP associated with income levels.

technologies may vary across income levels within a country in different ways than they vary across countries. We can discipline the latter using the time use data, but not the former.

[46] More genereally, these calculations suggest that the distribution of search willingness and efficiency will affect the prices paid in a market.

D. Biased Technology Gap

The column titled *Benchmark* in Table 11 reports the model's properties when the relative productivity of the tradable sector rises with income as in the Jorgensen-Stiroh data. With a biased productivity gap across countries, the model accounts for 51 percent of the PPP-income relation. The traditional HBS effect accounts for 28 percent of the price-income relationship, while pricing-to-market accounts for 23 percent. Pricing-to-market is similar to the balanced productivity case.

For comparison, in the column titled *HBS standard model*, the HBS model generates only 27 percent of the price-income relationship. Thus, given the observed productivity bias in the time-series data, the HBS model alone explains little of the PPP-income relationship. Our benchmark pricing-to-market model exhibits a slightly stronger HBS effect because it generates larger wage differences than income difference across countries. This occurs because higher market productivity raises consumption, which leads to more shopping time and less market labor. As income rises, the ratio of shopping to work time rises faster in our model than the point estimate in the data (0.968 vs. 0.330). Both models generate movements in the relative price of non-traded to traded goods with income that are only slightly higher than the data.

E. Sensitivity

We now consider the sensitivity of our results to the share of tradables, productivity bias, labor share of income, and shopping technology.[47] Except where noted, the model is parameterized as in the benchmark case of the biased productivity gap.

[47] We also explored varying the elasticity of substitution between tradables and non-tradables as well as the elasticity of substitution across tradables. Varying these elasticities primarily affected estimates of $\varepsilon_{N/T}$, but had very little impact on ε_{LOP} or ε_w. We also explored changing the level of shopping time per purchase equally across countries and this had a very minor impact on our estimates.

Share of Tradables

Figure 3 plots the relationship between the share of tradables and both pricing-to-market and violations of PPP with balanced productivity. The effect of varying the share of tradables is minor. This is because pricing-to-market affects all goods, traded and non-traded, in the same way. Thus, the tradables share only affects pricing-to-market through its influence on the terms of trade and in turn the relative wage. However, with highly substitutable goods this effect is small.

When the productivity gap is biased, the tradables share has a large effect on the size of violations from PPP. With a smaller share of tradables, non-traded goods receive a larger weight in prices. In Table 11, the column titled *Low Trade* reports the results of the benchmark model with a trade share of 20.9 percent. This is the necessary tradables share for the benchmark model to generate the same violations from PPP as in the data. This lower share of tradables slightly weakens the amount of pricing-to-market in the model.

For comparison, we also include the size of violations from PPP in the standard model with a low tradables share of $\alpha = 0.037$ in the column titled *HBS Low Trade*. This is the level of trade consistent with the aggregate price-income relationship in the standard HBS model. This tradables share generates trade flows that are only about 5 percent of those in the data and requires larger differences in the relative price of tradables to non-tradables than in the PWT data (0.50 vs. 0.41).

Biased Productivity

Figure 4 plots our measures of elasticities against the extent of comparative advantage in non-tradables (i.e., the ratio of relative non-tradable productivities to relative tradable productivities), which we denote as $g_{N/T} = \ln\left(a_1^{NT}/a_2^{NT}\right)/\ln\left(a_1^{T}/a_2^{T}\right)$. When $g_{N/T} = 0$, technological differences are completely concentrated in the tradables sector. When $g_{N/T} = 1$, there is no relative bias across sectors in technology levels. For comparison, the elasticity of deviations from PPP in a model without price discrimination is also reported as ε_{PPP_STD}.

From Figure 4 we see that the violations from PPP are decreasing in $g_{N/T}$, while pricing-to-market is increasing in $g_{N/T}$. To understand these different results, first note that in the model without price discrimination, ε_{PPP_STD} is decreasing in $g_{N/T}$ because the relative price of non-tradables is decreasing as the productivity gap diminishes.

Two factors influence the relationship between pricing-to-market and the productivity gap. First, firms face a lower bound on price in their pricing-to-market decision, since they will never charge below marginal cost. Thus, pricing-to-market is somewhat non-linear. Among relatively high-wage countries, firms will vary prices with their customers' wages, but among relatively low-wage destinations, markups are already quite low, so firms have very little ability to vary their price with the destination wage. Second, with a biased productivity gap, relative wage differences are much larger than relative income differences. This is because relative wages are determined primarily by the productivity difference in tradables, while relative income differences are based on productivity in both sectors. Taken together, these two features imply that a biased productivity gap leads to greater pricing-to-market among high-income locations and lower pricing-to-market among low-income locations. Given the world distribution of income, the reduced pricing-to-market to low-income locations has a stronger effect on the estimate of pricing-to-market.

From Figure 4, we see that for the model without pricing-to-market to account for the violations from PPP, the productivity gap in tradables must be 25 times the productivity gap in non-tradables, or about 15 times larger than in the U.S. time series data.

Labor Share

We now consider the effect of the labor share on the model's predictions. In Table 11, the column titled *Low Labor* reports the results of the model with a labor share of 50 percent. In this case, there are larger violations of PPP and these are entirely due to an increase in pricing-to-market. The lower labor share leads to larger markups and gives firms more room to price-to-market.

This is particularly important for pricing to low-income countries since firms will never price below marginal cost. With higher markups, the model now accounts for almost two-thirds (0.147/0.235) of the tradable price-income relationship and 75 percent (0.151/0.209) of the tradable price-wage relationship. Obviously, if we increase the labor share, and lower the average markup, we will weaken the price-income relation in the model.

Search Time

Our model relies on relative productivity differences in shopping to be smaller than in market production. As we have seen already, time-use surveys provide evidence of this, but not to the extent we have assumed. To make the model consistent with the time-use data, we allow the amount purchased per shopping trip to vary across countries with tradable productivity. We assume consumers in country j can purchase κ_j units per shopping trip, and let differences in κ_j be proportional to the differences in the tradable technology.[48] To match the elasticity of shopping to work time, we set the productivity gap in shopping to be 53 percent of the productivity gap in tradables, so that lower income countries continue to have a comparative advantage in shopping.

The column titled *Variable Shopping* reports the results of this modification. The estimates are quite similar to our benchmark for two reasons. First, because there is less substitution of work for shopping, we match the income distribution with smaller wage differences. This weakens the HBS effect. Second, the increased shopping time of lower income countries means that differences in search costs are smaller for a given difference in wages compared to the benchmark model. This leads firms to do less pricing-to-market to richer countries, but also allows them to do more pricing-to-market to poorer countries.[49] The net effect is a higher estimate of ε_{LOP} that counteracts the

[48] As before, we normalize the units per purchase in our base country to be $\kappa = 1$.

[49] This result may seem perverse but is largely due to estimating a linear model on non-linear data. If we plotted the distribution of prices against income in the benchmark model and the variable shopping model, we would find that for each income level the variable search model generates higher prices. However, the variable search model generates an almost linear relationship between income and prices, while the benchmark model generates a non-linear relationship.

lower HBS effect, leading to a very small change in the aggregate price-income relationship.

F. Relative Prices and Relative Wages

Pricing-to-market in the model is driven by the opportunity cost of time measured by wages and not income per capita. As the model abstracts from important determinants of income per capita such as population growth, labor market participation, and capital accumulation, focusing on relative wages and prices is a more direct test of the model. Rather than match the distribution of income per capita, we recalibrate technology to match the exact distribution of relative nominal wages in the sample of countries for which we have wage data. Figure 5 plots the relationship between relative price levels and wages from the model and the data.

In the data, the elasticity of price levels with respect to wages, which we denote ε_w^{PPP} is equal to 0.41. The model generates $\varepsilon_w^{PPP} = 0.30$ and thus can explain nearly 75 percent of the relationship between prices and wages. We find that pricing-to-market is the largest source of the price-wage relationship, since it accounts for two-thirds of this, while HBS effect accounts for one-third. The stronger pricing-to-market relationship that we find with wages is consistent with our empirical result, in which wages seem to drive the pricing-to-market.[50]

IV Conclusions

Using highly disaggregated data on U.S. exports at the border, we find strong empirical evidence that pricing-to-market accounts for a substantial amount of the long-run differences in tradable prices across countries. In turn, these tradable price differences are an important source of the deviations from absolute PPP, accounting for about 40 percent of the relation between aggregate price levels and income per capita in the data. This is in stark contrast to the conventional view that deviations from absolute PPP are solely due to differences in non-traded goods prices.

[50] Of course, it is also true that since wages are our main focus, this may be due to our rather simplistic modeling of income, which ignored leisure decisions, capital income, etc.

Our empirical work suggests that consumers in low-income countries are more price sensitive than consumers in high-income countries. We develop a model with this type of pricing-to-market based on international productivity differences and search frictions. Similar to HBS, our model relies on low-income countries having a comparative advantage in producing non-traded goods. Unlike HBS, these non-traded goods are shopping activities that affect all prices. Our model generates a role for local wages in the price-setting behavior of firms and is consistent with cross-country differences in shopping activities. The model is also consistent with two features of our data analysis suggesting an important role for search frictions. First, contrary to previous work, we find that wages have substantially more explanatory power for pricing-to-market than income per capita. Second, pricing-to-market appears strongest for those goods for which search frictions are likely to be most important, consumer goods and goods sold in decentralized transactions. This evidence is also consistent with the within-country evidence that consumers can use search to lower their average purchase price. Our model thus provides a simple, unified theory of within and across country price dispersion.

Comparing the results of our quantitative analysis and empirical work offers a way to quantify the magnitude of different sources of international price differences in tradables. Assuming our model captures all the reasons for pricing-to-market by income, we find that pricing-to-market (0.109/0.310=35 percent), quality differences ((0.235-0.109)/0.310=41 percent), and downstream non-traded inputs ((0.310-0.235)/0.310=24 percent) contribute nearly equally in the differences in tradable prices by income. Naturally, other factors may contribute to the pricing relation we have found, so further empirical work is necessary. Nevertheless, ours is, to the best of our knowledge, the first study to document an important empirical and theoretical role for tradable prices in absolute PPP. It is typically assumed in theoretical and empirical work that the LOP holds for tradables. We have shown that such an assumption drastically overstates the differences in productivity across sectors across countries and matters for understanding the source of income differences as well.

References

[1] Aguiar, Mark, and Erik Hurst. 2007. "Lifecycle Prices and Production." American Economic Review, 97(5): 1533-1559.

[2] Alessandria, George. 2009. "Consumer Search, Price Dispersion and International Relative Price Variability." International Economic Review, 50(3).

[3] Alessandria, George. 2004. "International Deviations from the Law of One Price: The Role of Search Frictions and Market Share." International Economic Review, 45(4): 1263-91.

[4] Alessandria, George, and Joseph Kaboski. 2007. "Pricing-to-Market and the Failure of Absolute PPP." FRB of Philadelphia Working Paper No. 07-29.

[5] Alessandria, George, Joseph Kaboski, and Virgiliu Midrigan. Forthcoming. "Inventories, Lumpy Trade and Large Devaluations." American Economic Review.

[6] Asea, Patrick, and Enrique Mendoza. 1994. "The Balassa-Samuelson Model: A General Equilibrium Appraisal." Review of International Economics, 2(3): 244-67.

[7] Atkeson, Andrew, and Ariel Burstein. 2008. "Pricing-to-Market, Trade Costs, and International Relative Prices." American Economic Review, 98(5): 1998-2031.

[8] Balassa, Bela. 1964. "The Purchasing Power Parity Doctrine: A Reappraisal." Journal of Political Economy, 72: 584-96.

[9] Bergstrand, Jeffrey H. 1991. "Structural Determinants of Real Exchange Rates and National Price Levels: Some Empirical Evidence." American Economic Review, 81(1): 325-34.

[10] Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." Econometrica, 63(4): 841-890.

[11] Belsley, David A. 1969. Industry Production Behavior: The Order-Stock Distinction. Amsterdam: North-Holland Pub. Co.

[12] Besedes, Tibor, and Thomas J. Prusa. 2006. "Product Differentiation and Duration of U.S. Import Trade?" Journal of International Economics, 70(2): 339-58.

[13] Bhagwati, Jagdish N. 1984. "Why Are Services Cheaper in Poor Countries?" Economic Journal, 94(2): 79-86.

[14] Bils, Mark, and James Kahn. 2000. "What Inventory Behavior Tells Us about Business Cycles." American Economic Review, 90(3): 458-481.

[15] Burdett, Kenneth, and Kenneth L. Judd. 1983. "Equilibrium Price Dispersion." Econometrica, 51(4): 955-70.

[16] Canzoneri, Matthew, Robert Cumby, and Behzad Diba. 1999. "Relative Labor Productivity and the Real Exchange Rate in the Long Run: Evidence for a Panel of OECD Countries." Journal of International Economics, 47: 245-266.

[17] Chari, V.V., Patrick J. Kehoe, and Ellen McGrattan. 2002. "Can Sticky Price Models Generate Volatile and Persistent Real Exchange Rates?" Review of Economic Studies, 69(3): 533-63.

[18] Cooley, Thomas F., and Edward C. Prescott. 1996. "Economic Growth and Business Cycles." In Frontiers of Business Cycle Research, ed. Thomas Cooley, 1-38. Princeton: Princeton University Press.

[19] Dornbusch, Rudiger. 1987. "Exchange Rates and Prices." American Economic Review, 77: 93-106.

[20] Dornbusch, Rudiger. 1988. "Purchasing Power Parity." In The New Palgrave Dictionary of Economics, ed. John Eatwell, Murray Milgate, and Peter Newman, 1075-85. London: MacMillan.

[21] Dotsey, Michael, Robert King, and Alexander L. Wolman. 1999. "State-dependent Pricing and the General Equilibrium Dynamics of Money and Output." Quarterly Journal of Economics, 114(2): 655-690.

[22] Eaton, Jonathan, and Samuel Kortum. 2001. "Trade in Capital Goods." European Economic Review, 45(7): 1195-235.

[23] Engel, Charles. 1999. "Accounting for U.S. Real Exchange Rates." Journal of Political Economy, 107(3): 507-38.

[24] Feenstra, Robert. 1989. "Symmetric Pass-Through of Tariffs and Exchange Rates Under Imperfect Competition: An Empirical Test." Journal of International Economics, 27(2): 25-45.

[25] Feenstra, Robert, John Romalis, and Peter Schott. 2002. "U.S. Imports, Exports and Tariff Data, 1989-2001." NBER Working Paper 9387.

[26] Goldberg, Pinelopi. 1995. "Product Differentiation and Oligopoly in International Markets: The Case of U.S. Automobile Industry." Econometrica, 63(4): 891-951.

[27] Hallak, Juan C. 2006. "Product Quality and the Direction of Trade." Journal of International Economics, 68(1): 238-265.

[28] Harrod, Roy F. 1933. International Economics. London: James Nisbet and Cambridge University Press.

[29] Herrendorf, Berthold, and Akos Valentinyi. 2006. "Which Sectors Make the Poor Countries so Unproductive?" Society for Economic Dynamics 2006 Meeting Paper 204.

[30] Hsieh, Chang-Tai, and Peter J. Klenow. 2007. "Relative Prices and Relative Prosperity." American Economic Review, 97: 562-585.

[31] Hummels, David L., and Alexandre Skiba. 2004. "Shipping the Good Apples Out? An Empirical Confirmation of the Alchian-Allen Conjecture." Journal of Political Economy, 112(6): 1384-1402.

[32] Isard, Peter. 1977. "How Far Can We Push the 'Law of One Price'?" American Economic Review, 67(5): 942-48.

[33] Jorgenson, Dale W., and Kevin J. Stiroh. 2000. "U.S. Economic Growth at the Industry Level." American Economic Review, 90(2): 161-7.

[34] Knetter, Michael. 1993. "International Comparisons of Pricing-to-Market Behavior." American Economic Review, 83(3): 473-86.

[35] Kravis, Irving and Robert E. Lipsey. 1983. Toward an Explanation of National Price Levels. Princeton, NJ: Princeton University International Finance Section.

[36] Krugman, Paul. 1987. "Pricing to Market when the Exchange Rate Changes." In Real Financial Linkages Among Open Economies, ed. Sven Arndt and J. David Richardson, 49-70. Cambridge, MA: MIT Press.

[37] Linder, Staffan B. 1961. An Essay on Trade and Transformation. New York: Wiley.

[38] McKenzie, David, and Ernesto Schargrodsky. 2005. "Buying Less, but Shopping More: Changes in Consumption Patterns During a Crisis." BREAD Working Paper No. 092.

[39] Mendoza, Enrique. 1995. "The Terms of Trade, the Real Exchange Rate, and Economic Fluctuations." International Economic Review, 36(1): 101-37.

[40] Neary, Peter. 1988. "Determinants of the Equilibrium Real Exchange Rate." American Economic Review, 78(1): 210-215.

[41] Nevo, Avi. 2001. "Measuring Market Power in the Ready-to-Eat Cereal Industry." Econometrica, 69(02): 307-342.

[42] Obstfeld, Maurice, and Kenneth Rogoff. 1996. Foundations of International Macroeconomics. Cambridge, MA: MIT Press.

[43] Ostry, Jonathan, and Carmen Reinhart. 1991. "Private Saving and Terms of Trade Shocks: Evidence from Developing Countries." IMF Working Paper No. 91/100.

[44] Rauch, James E. 1999. "Networks versus markets in international trade." Journal of International Economics, 48: 7-35.

[45] Rauch, James E. 2001. "Business and Social Networks in International Trade." Journal of Economic Literature, 39: 1177-1203.

[46] Rauch, James E., and Alessandra Casella. 2003. "Overcoming Informational Barriers to International Resource Allocation: Prices and Ties." Economic Journal, 113: 21-42.

[47] Rauch, James E., and Joel Watson. 2003. "Starting Small in an Unfamiliar Environment." International Journal of Industrial Organization, 21: 1021-1042.

[48] Rogoff, Kenneth. 1996. "The Purchasing Power Parity Puzzle." Journal of Economic Literature, 34: 647-668.

[49] Samuelson, Paul. 1964. "Theoretical Notes on Trade Problems." Review of Economics and Statistics, 46: 145-154.

[50] Schott, Peter. 2004. "Across-Product versus Within-Product Specialization in International Trade." Quarterly Journal of Economics, 119: 647-668.

[51] Simonovska, Ina. 2008. "Income Differences and Prices of Tradables." mimeo.

[52] Stockman, Alan, and Linda Tesar. 1995. "Tastes and Technology in a Two-Country Model of the Business Cycle: Explaining International Comovements." American Economic Review, 85(1): 168-185.

Appendix A - Random Sample of 10 Goods

Consumer/Automotive Goods

- Alarm clocks, not battery or AC powered
- Men's and boys's weaters of cotton, knitted or crocheted containing ge 36 percent by weight of flax fiber
- Pocket lighters, gas fueled, refillable
- Table or kitchen glassware other than drinking glasses, having a linear coefficient of expnsn nov 5X0-6/Kelvin in temperature range of 0 to 300 deg C
- Washing machines, except coin operated, fully automatic, dry linine capacity not exceeding 10 kg, household or laundry type

Non-Consumer Goods (Agriculture, Industrial, and Equipment)

- 4,4'-Isopropylidendicyclohexanol; & Mixtures cont less than 90% by weight of stereoisomers of 2-Isopropyl-5-Methylcylohexanol, but not more than.
- Grinders, polishers and sanders, suitable for metal working, rotary type (inc combined rotary-percussion) pneumatic tools for working in the hand
- Monolithic I/C's, digital, silicon, (MOS), volatile memory, static read-write random access (SRAM) over 300,000 bits
- Peanuts, blanched
- Synthetic filament yarn except sewing thread, not for retial sale, single, multifilament, with a twist of GE turns per M of polyethylene, propylene

Figure 1: Consumption Price Levels and Real GDP per Capita

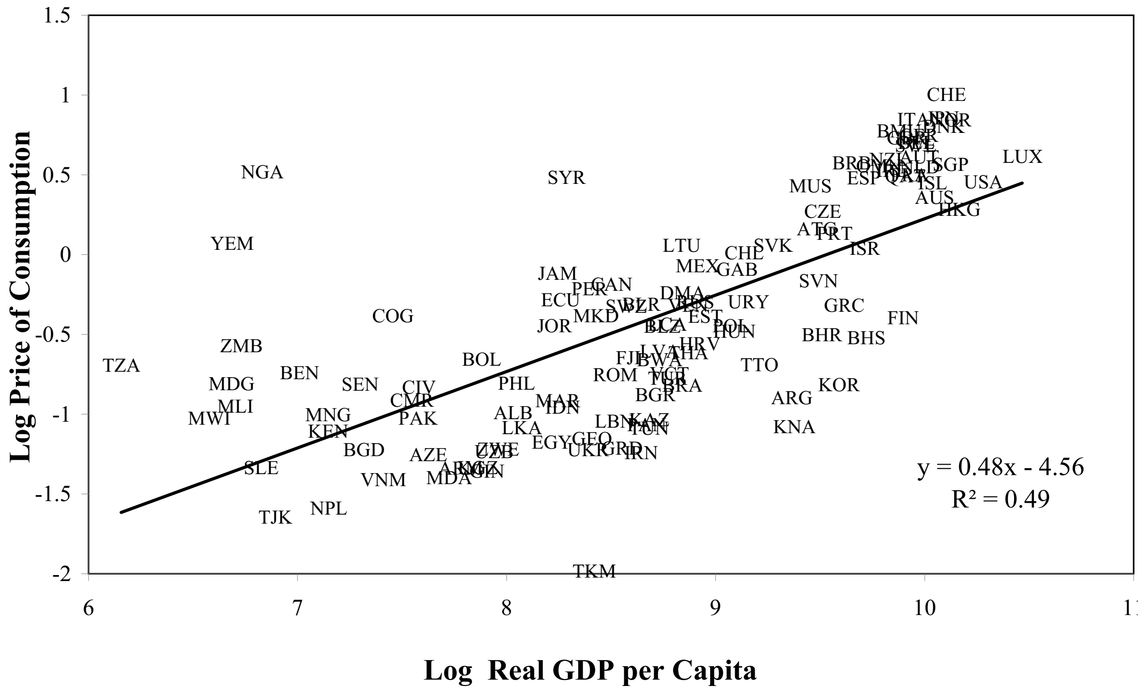

(Source: Penn World Tables 6.1, ICP 1996 Benchmark Price Data)

Figure 2: Tradable Consumption Prices and Real GDP per Capita

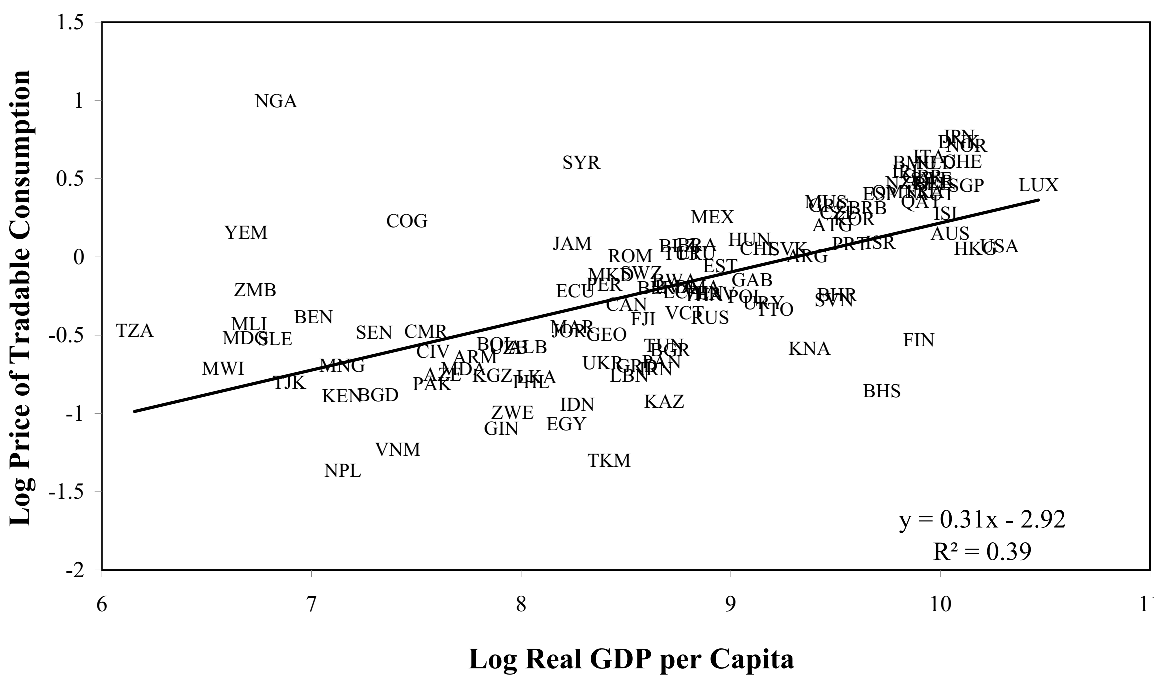

(Source: Penn World Tables 6.1, ICP 1996 Benchmark Price Data)

Figure 3: Price Elasticities and Tradable Share

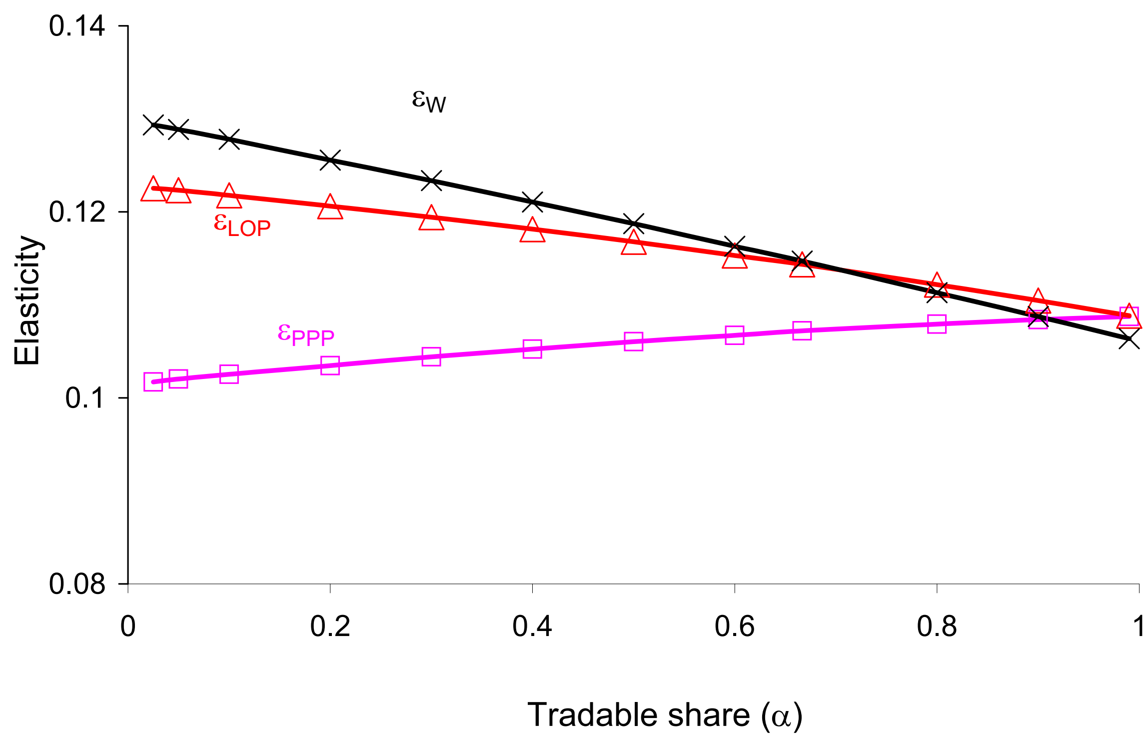

Figure 4: Price Elasticities and Productivity Gap

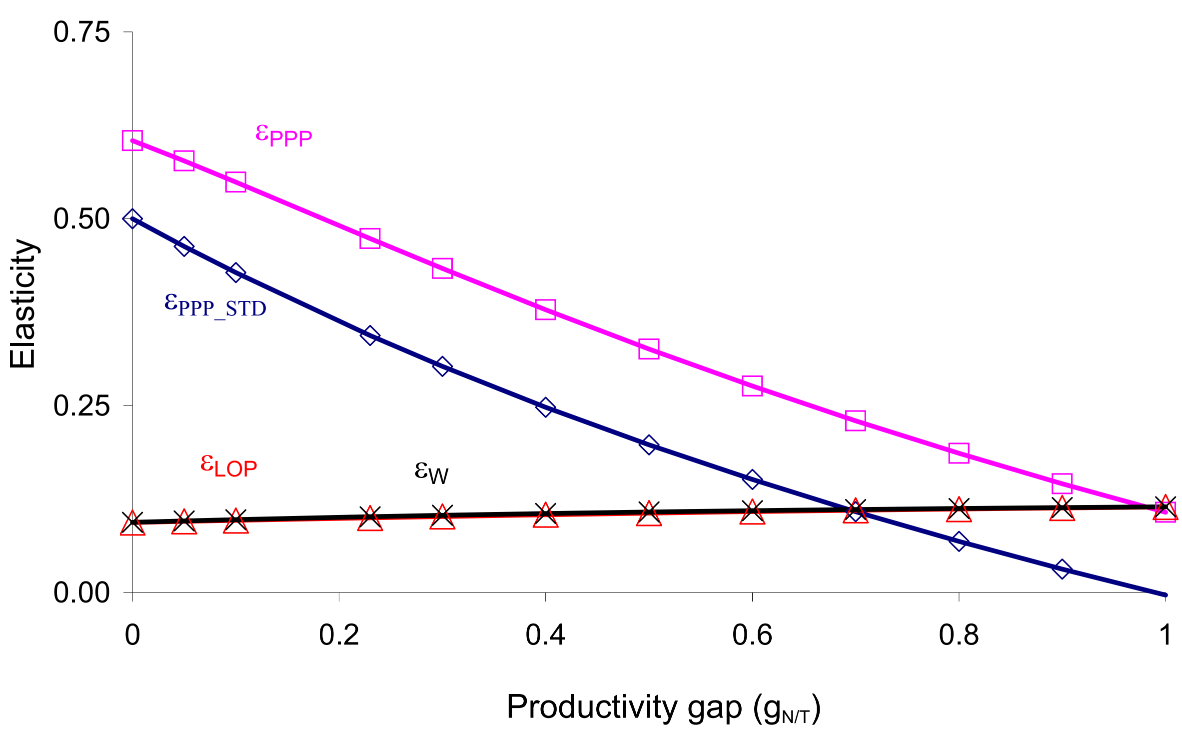

Figure 5: Relative Prices and Wages

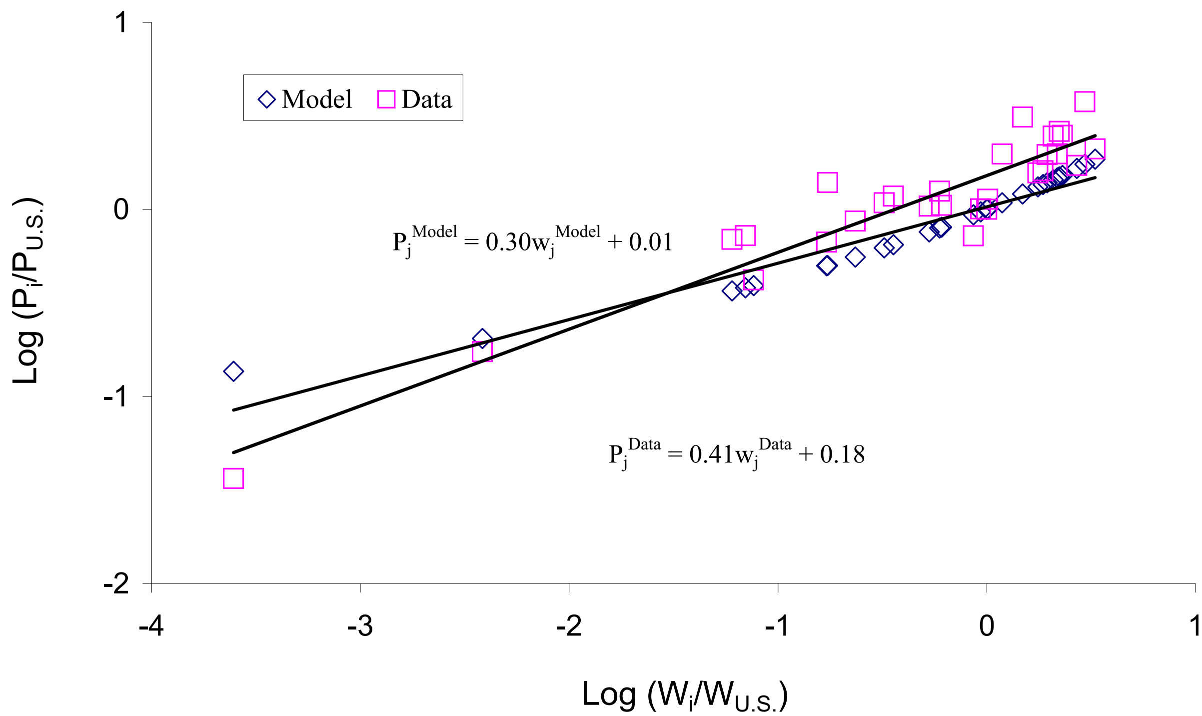

Table 1: Coefficients from Commodity-Year Fixed-effects Regressions of Log Unit Values on Log real GDP per capita and/or Log Wages

(tstatistics in parentheses)

Coefficient	GDP per Capita only	Wage only	Both together
Log GDP per capita	0.235 (4.9)	-	0.041 (1.4)
Log Wage	-	0.209 (7.1)	0.188 (8.1)
Observations	240245	240245	240245
Groups	5613	5613	5613

t-statistics are based on country-year clustered White robust standard errors.

Table 2: PTM and "Quality" Variation using Import Prices

(tstatistics in parentheses)

Variable	All	Small quality differences	Large quality differences	Smallest quality differences	Interacting quality differences
Log Wage	0.159	0.128	0.183	0.107	0.084
	(5.7)	(6.2)	(5.2)	(5.4)	(4.7)
Observations	704121	296186	407935	44552	704121
Groups	49513	24757	24756	4951	49513
Mean quality	1.44	0.78	1.92	0.29	1.44

Based on those HS codes that are both imported and exported in the same year. For each commodity (a year hs code) quality variation is measured as the standard deviation of ln import unit values from all sources. For columns 2 and 3, goods with small quality differences have a standard deviation of prices of less than 1.21 and goods with large quality differences have a standard deviation greater than 1.21. For column 4, goods with the smallest quality differences have a standard deviation of unit values of less than 0.46. We run the regression from equation 1 on these seperate groups with good fixed effects. The final column reports the results of the regression in equation 2 that includes an interaction of quality and relative income. The row Mean quality reports the observation weighted mean quality of each sample. t-statistics are based on country-year clustered White robust standard errors.

Table 3 : Export Prices and Wages Using Monthly Data from 1998 and 1999

(tstatistics in parentheses)

	Monthly-port data				Annual port		Monthly-hs	Annual-hs	
Variable\grouping	port-hs-month	port-hs-year	hs-month	hs-yr	port-hs-yr	hs-yr	hs-month	hs-yr	hs-yr, no additional controls
Log Wage	0.096	0.087	0.085	0.082	0.085	0.081	0.148	0.112	0.184
	(2.8)	(2.6)	(3.1)	(3.1)	(3.4)	(5.8)	(12.6)	(4.2)	(2.8)
Observations	3532393	3532393	3532393	3532393	848432	848432	1322053	217707	217707
Groups	1753642	289895	170184	15083	289895	15083	87954	15083	15083

Notes: Includes separate controls for share of trade (value) shipped by air and ln number of shipments; t-statistics are based on country-year clustered White robust standard errors.

Table 4: Effect of Quality Aggregation on Coefficients from Regressions of Log Prices on Log Real Wages

(tstatistics in parentheses)

	Number of Commodity Groupings	Number of Observations	Coefficient
Panel A: All Commodities			
Combined at the 9-digit level			
Individual Commodities	855	74264	0.184
			(6.3)
Aggregated Commodities	338	48552	0.158
			(5.0)
Combined at the 7-digit level			
Individual Commodities	4186	459542	0.150
			(5.9)
Aggregated Commodities	1155	209972	0.115
			(3.8)
Combined at the 5-digit level			
Individual Commodities	7358	861030	0.164
			(6.1)
Aggregated Commodities	1490	320013	0.134
			(4.4)
Panel B: Commodities Grouped by Identifiable Quality Differences			
Individual Commodities	7593	88340	0.151
			(7.6)
Aggregated Commodities	3492	51152	0.126
			(6.2)

Notes: For Panel B there are 1112 separate commodities that are grouped in 455 groups. t-statistics are based on country-year clustered White robust standard errors.

Table 5: PTM By Enduse Category and Rauch Classification

(tstatistics in parentheses)

	All Goods	Foods, feeds, and beverages (EU=0)	Industrial supplies and materials (EU=1)	Capital goods, except automotive (EU=2)	Auto, vehicles, parts and engines (EU=3)	Consumer goods (EU=4)
Log Wage	0.164	0.091	0.176	0.143	0.11	0.218
	(5.9)	(5.2)	(5.7)	(4.3)	(3.6)	(7.4)
Observations	1125186	104774	464314	308485	24715	215530
By Rauch Classification and End-use						
Rauch = "w" - "goods traded on an organized exchange (homogeneous goods)"						
Log Wage	0.074	0.042	0.110			
	(3.2)	(2.3)	(3.5)			
Observations	38865	19656	19209			
Rauch = "r" - "reference priced goods"						
Log Wage	0.146	0.107	0.158	0.378		0.152
	(5.3)	(5.9)	(5.0)	(4.5)		(6.0)
Observations	204681	48810	148388	790		6693
Rauch = "n" - "differentiated goods"						
Log Wage	0.163	0.096	0.176	0.123	0.095	0.221
	(5.6)	(5.8)	(5.1)	(4.1)	(2.8)	(7.7)
Observations	547594	29291	159797	185027	15148	153673

Notes: t-statistics are based on country-year clustered White robust standard errors. Rauch classification based on conservative classification scheme.

Table 6: Role of Inventory Holdings for Pricing-to-Market

(tstatistics in parentheses)

	To Order (Belsley)	To Stock (Belsley)	To Order (Bils & Kahn)	To Stock (Bils & Kahn)	Final good share of inventory (bottom quartile)	Final good share of inventory (top quartile)
Log Wage	0.160	0.185	0.153	0.207	0.146	0.191
	(5.0)	(7.6)	(5.0)	(8.1)	(4.9)	(8.1)
Observations	651595	318129	700885	268839	266007	237294
Final good inventory share	0.354	0.518	0.371	0.503	0.240	0.600

Belsley (1969) classifies 6 2-digit SIC industries as more likely to produce to stock. The sectors are: food and kindred products, tobacco, apparel and other textile products, chemicals and allied products, rubber, and plastic products. Bils and Kahn (2000) modify Belsley by shifting food into to order and lumber to stock. Inventories are measured at the 6-digit NAICS level (except for 5 industries at the 5-digit level) using the 1997 Census of Manufacturers. The share of final inventories measures the finished inventories divided by total inventory holdings (where inventory measures are an average of beginning- and end-of-year holdings).

Table 7: Pricing-to-market and Repeated Transactions

(tstatistics in parentheses)

Variable	Whole yr (99)	99Q4	99Q4
Log Wage	0.229	0.215	0.194
	(2.5)	(2.4)	(2.7)
1 previous transaction			-0.049
			(2.0)
2 previous transactions			-0.079
			(3.0)
3 previous transactions			-0.345
			(6.6)
Observations	298460	75609	75609

Notes: A previous transaction is a quarter with a shipment to a country. Dummies for the number of previous transactions are based on the number of quarters in the first 3 quarters of the year with transactions. t-statistics are based on country-year clustered White robust standard errors. Rauch classification based on conservative classification scheme.

Table 8: Coefficient from Regression of Log Time Use on Log GDP per capita
(t-statistics in parentheses)

Group	Observations	Shop/Work Time	Work Time*	Shop Time**
EHTUS	15	0.337	-0.139	0.198
		(2.8)	(-2.0)	(2.8)
MTUS	48	0.321	-0.112	0.208
		(4.3)	(-2.7)	(3.1)

EHTUS countries: Belgium, Estonia, Finland, France, Germany, Hungary, Italy, Latvia, Lithuania, Norway, Poland, Slovenia, Spain, Sweden, U.K. and sample is 20 to 74 year olds. MTUS countries: Canada (71, 81, 86, 92, 98), Denmark (64, 87) France (65, 74, 98), Netherlands (75, 80, 85, 90, 95, 00), Norway (71, 81, 90, 00), U.K. (61, 75,83, 87, 95,00), USA (65, 75, 85, 92, 98, 03), Hungary (65, 77), West Germany (65), Poland (65), Belgium (65), Bulgaria (88), Czechslovakia (65), East Germany (65), Yugoslavia (65), Italy (80, 89), Australia (74), Israel (92), Germany (92), Austria (92), S. Africa (00), Slovenia (00) and the sample is 20 to 59 year olds.

* EHTUS work time is measured as paid work in primary and secondary employment. MTUS work time is measured as paid work in first and second job plus paid work at home.

** Shop time is measured as time shopping and receiving personal services plus time travelling to shopping.

Table 9: U.S. Sectoral Productivity Growth (1958-96)

	Average Growth		Weighted Avg* Growth	
	TFP	Labor Productivity	TFP	Labor Productivity
Tradables	0.56	1.85	0.67	2.07
Non-Tradables	0.27	1.60	0.26	1.41
Ratio ($g_{N/T}$)	0.48	0.87	0.38	0.68

Tradables include: Agriculture; Metal Mining; Coal Mining; Petroleum and Gas; Nonmetallic Mining; Food Products; Tobacco Products; Textile Mill Products; Apparel and Textile; Lumber; Furniture; Paper Products; Printing and Publishing; Chemical Products; Petroleum Refining; Rubber and Plastics; Leather Products; Primary Metals; Fabricated Metals; Industrial Machinery and Equipment; Electronic and Electric Equipment; Motor Vehicles; Instruments; Miscellaneous Manufacturing; Other Transportation Equipment; Stone, Clay and Glass.

Non-tradables include: Construction; Transport and Warehouse; Communications; Electric Utilities; Gas Utilities; Trade; FIRE; Services.

*The weighted measure weights productivity growth in each sector by its average annual share of value added in either the tradable or non-tradable sector. The productivity data is reported in Jorgenson and Stiroh (2000). Total Factor Productivity Growth (TFP) is measured as a residual using materials, capital stocks, and labor used plus their expenditure shares.

Table 10: Parameter Values and Measurement

Parameters	
Benchmark Model	
Technology	$\bar{a}=1/4$, $q=0.727$, $\ln(a_i^{NT}/\bar{a})/\ln(a_i^{T}/\bar{a}) = 2/3$
Preferences	$\alpha=2/3, \rho=0.99$
Variations	
Balanced	$\ln(a_i^{NT}/\bar{a})/\ln(a_i^{T}/\bar{a}) = 0$
Low Trade Share	$\alpha=0.209$
Low Labor Share	$q=0.8$, $\bar{a}=1/4$
Vary Shop	$(\kappa_i^{NT}/\ \bar{\kappa}_i)/(a_i^{T}/\bar{a}) = 0.53$
HBS Model	$q=0.00005$, $\bar{a}=1/4$
Standard	$\alpha=2/3, \rho=0.99$
Low Trade	$\alpha=0.037$
Measurement	
Nominal Income	$Y_j=P_{j1}(R_{j1})N_{j1}+P_{j2}(R_{j2})N_{j2}+P_{j3}(R_{j3})N_{j3}$
Aggregate Prices	$P_j=(P^T_j/\alpha)^{\alpha}(P_{3j}/(1-\alpha))^{1-\alpha}$, $P_j^T = ((P_{1j})^{(\rho-1)/\rho}+(P_{2j})^{(\rho-1)/\rho})^{1-1/\rho}$
Empirics	$\ln P_j = \varepsilon_{PPP}\ln y_j+e_j$
	$\ln LOP_j = \varepsilon_{LOP}\ln y_j+e_j$
	$\ln LOP_j = \varepsilon_{w}\ln w_j+e_j$
	$\ln P_j^{NT}/P_j^{T} = \varepsilon_{N/T}\ln y_j+e_j$

Table 11: Model Results

			Variations on Biased Productivity Economies**					
			PTM Model				HBS Model	
		Balanced		Low Labor	Variable	Low Trade		Low Trade
A. Elasticity*	Data	PTM	Benchmark	share = 1/2	Shopping	α=0.209	Standard	α=0.037
ε_{PPP}	0.480	0.107	0.245	0.283	0.232	0.480	0.128	0.430
ε_{LOP}	0.235	0.114	0.109	0.147	0.117	0.111	0.000	0.000
ε_{w}	0.209	0.115	0.111	0.151	0.115	0.115	0.000	0.000
$\varepsilon_{N/T}$	0.408	-0.022	0.406	0.410	0.345	0.467	0.384	0.498
$\varepsilon_{shop/work}$	0.330	1.041	0.968	0.962	0.332	1.001	0.988	1.042
B. Accounting for violations from PPP								
Fraction of ε_{PPP}		22%	51%	59%	48%	100%	27%	90%
Fraction from PTM		24%	23%	31%	24%	31%	0%	0%
Fraction from HBS		-1%	28%	28%	24%	69%	27%	90%

*ε_{PPP} and $\varepsilon_{N/T}$ are based on the whole sample of 115 PWT Benchmark countries while ε_{w} and ε_{LOP} are based the 28 benchmark countries for which the BLS provides wage data.

**The variations of the Benchmark economy all include a biased productivity gap. In the HBS economies there is no pricing-to-market but consumers do shop for goods. Low labor share is the Benchmark economy with labor share of income of 1/2. The Variable Shopping economy is one in which the shopping technology improves along with the tradable technology. The Low Trade Share economy is the Benchmark economy with a lower tradable share of 0.209. The HBS Low Trade is the HBS model with a low trade share of 0.037.